

Knowledge discovery from data sets under tenable assumptions

Marco Zaffalon



Istituto Dalle Molle di Studi sull'Intelligenza Artificiale
SWITZERLAND
<http://www.idsia.ch/~zaffalon>
zaffalon@idsia.ch

School on Imprecise Probabilities, Lugano (CH), 27–31 July 2004

2

Introduction to Knowledge discovery from data sets (KDD)

Prototypical framework

- Data (d) come in a table
- $c_1 \quad a_{11} \quad \dots \quad a_{1j} \quad \dots \quad a_{1m}$
- $c_i \quad a_{i1} \quad \dots \quad a_{ij} \quad \dots \quad a_{im}$
- $c_N \quad a_{N1} \quad \dots \quad a_{Nj} \quad \dots \quad a_{Nm}$
- $c_i \in \mathcal{C} \quad (i = 1, \dots, N)$
- \mathcal{C} is a finite set of *classes*
- $a_{ij} \in \mathcal{A}_j \quad (i = 1, \dots, N; j = 1, \dots, m)$
- \mathcal{A}_j 's are finite sets of *attribute values*
- A row of the table is called *observation*
- The data in the table are also called *sample*
- N , the number of units, is also called *sample size*

3

KDD

- Learning models from data (alone)
- Using models to make “predictions” about new data
- Discipline related to AI and Statistics
- Emphasis on computer-intensive methods

4

(Pattern) Classification

- Use data to learn a function $f : \mathcal{A}_1 \times \cdots \times \mathcal{A}_m \rightarrow \mathcal{C}$
 - f is "learned by examples"
 - f is called *classifier*
 - Learning = inferential approach
- Use the classifier to predict the *unknown* class of a *new* observation
 - $(a_1, \dots, a_m) \xrightarrow{f} c$
- Why being interested in classification?

5

Some applications

- Hand-written character recognition
- Observation = Image
 - Classes = {a, b, ..., z}
- Face recognition
- Medical diagnosis
 - List of Symptoms
 - Disease
- ... Fraud detection, network intrusion, gene expression profiling, ...
 - ... many others: very general paradigm

6

Common assumption

- The process underlying the data is *multinomial*
 - Observations are generated in independent and identically distributed way
 - C is called *class variable*, A_j ($j = 1, \dots, m$) *attribute variable*
- Notation
 - $\mathcal{X} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_m$, with generic element $x = (a_1, \dots, a_m)$
 - $\mathcal{D} = \mathcal{C} \times \mathcal{X}$, with generic element $d = (c, x)$
- The *physical probability* (or *chance*) of d is θ_d ($\sum_{d \in \mathcal{D}} \theta_d = 1$)
 - Let θ be the vector of chances
 - θ is usually unknown
- The probability $p(d|\theta)$ of the observed data is then $\prod_{d \in \mathcal{D}} \theta_d^{n_d}$
 - n_d is the number of times the instance d appears in the data set \mathcal{D}
 - As a function of θ , $p(\mathbf{d}|\theta)$ is called the *likelihood function*

7

Bayesian approach to classification

- Model your state of knowledge about θ by a *prior* density $p(\theta)$
- Compute the *posterior* density $p(\theta|\mathbf{d})$ by Bayes' rule
 - $$p(\theta|\mathbf{d}) = \frac{p(\theta)p(\mathbf{d}|\theta)}{\int p(\theta)p(\mathbf{d}|\theta)d\theta}$$
- Focus on $p(c, x|\mathbf{d}) = E[\theta_{c,x}|\mathbf{d}]$
 - E denotes expectation w.r.t. $p(\theta|\mathbf{d})$
- Select $c^* = \operatorname{argmax}_{c \in \mathcal{C}} p(c|x, \mathbf{d}) = \operatorname{argmax}_{c \in \mathcal{C}} p(c, x|\mathbf{d})$

The multinomial case

- Traditional choice for $p(\theta)$ is Dirichlet(s, \mathbf{t}): $p(\theta) \propto \prod_{d \in \mathcal{D}} \theta_d^{s_{t_d}-1}$
 - s and $\mathbf{t} = (t_d)_{d \in \mathcal{D}}$ are hyperparameters
 - t_d = expectation of θ_d w.r.t. the prior
 - Informal* view (prior is conjugate)
 - $s > 0$ as the number of additional (fractional) *virtual* observations
 - t_d as the proportion of observations equal to d in the virtual sample
 - $0 < t_d < 1, \sum_{d \in \mathcal{D}} t_d = 1$
 - e.g., uniform prior: $s = |\mathcal{D}|, t_d = 1/|\mathcal{D}|$; Perks prior: $s = 1, t_d = 1/|\mathcal{D}|$
 - Posterior is Dirichlet($N+s, \mathbf{t}^*$): $p(\theta | \mathbf{d}) \propto \prod_{d \in \mathcal{D}} \theta_d^{n_{t_d}+s_{t_d}-1}$
 - $\mathbf{t}^* = (t_d^*)_{d \in \mathcal{D}}, t_d^* = (n_d + s_{t_d}) / (N + s)$
 - t_d^* = expectation of θ_d w.r.t. the posterior
 - The needed probability is then
- $$p(c, x | \mathbf{d}) = E[\theta_{c,x} | \mathbf{d}] = \frac{n_{c,x} + s_{t_c,x}}{N + s}$$

9

Some popular classifiers

10

The naive Bayes classifier (NBC)

- Unstructured models do not work well when \mathcal{D} large/small sample
- Problem of *overfitting*
 - The model memorizes the data rather than learning from them \Rightarrow bad predictions
- Structural assumption
 - Attribute variables are mutually independent given class variable, i.e.:
 - $\theta_{c,x} = \theta_c \prod_{j=1}^m \theta_{a_j|c}$ for all (c, x)
 - θ_c is the chance of $C = c$
 - $\theta_{a_j|c}$ is the chance of $A_j = a_j$ given $C = c$
- Simple but effective classifier, and very popular
 - Independence assumption not critical for classification tasks

11

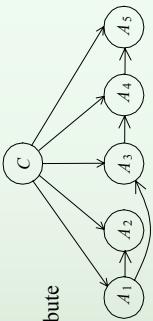
Bayesian approach to NBC classification

- Independence assumption makes the likelihood factorize
- $$\prod_{(c,x) \in \mathcal{D}} \theta_{c,x}^{n_{c,x}} = \prod_{c \in \mathcal{C}} \theta_c^{n_c} \prod_{j=1}^m \prod_{a_j \in \mathcal{A}_j} \theta_{a_j|c}^{n_{c,a_j}}$$
- Similarly for the prior, and the posterior is then proportional to
- $$\prod_{c \in \mathcal{C}} \theta_c^{n_c + s_{t_c}} \prod_{j=1}^m \prod_{a_j \in \mathcal{A}_j} \theta_{a_j|c}^{n_{c,a_j} + s_{t_{c,a_j}}}$$
- which is a product of Dirichlet densities, from which
- $$p(c, x | \mathbf{d}) = E[\theta_{c,x} | \mathbf{d}] = \frac{n_c + s_{t_c}}{N + s} \prod_{j=1}^m \frac{n_{c,a_j} + s_{t_{c,a_j}}}{n_c + s_{t_c}} = p(c | \mathbf{d}) \prod_{j=1}^m p(a_j | c, \mathbf{d})$$

12

Relaxing the independence assumption: the TAN model

- Tree augmented Bayes (TAN)
 - The dependence structure between attribute variables, given class variable, is a tree



- Learning problem
 - Learning the tree structure + probabilities
 - Mix of maximum likelihood and Bayesian techniques
- Classification

$$p(c, x | \mathbf{d}) = p(c | \mathbf{d}) \prod_{j=1}^m p(a_j | \pi_{A_j}, \mathbf{d})$$

13

Learning TAN structure

- Consider the conditional mutual information (MI)

$$\text{MI}(A_j, A_k | C) = \sum_{c \in C} \theta_c \sum_{(a_j, a_k) \in A_j \times A_k} \frac{\theta_{a_j, a_k | c}}{\theta_{a_j | c} \theta_{a_k | c}}$$
- Procedure
 - Set up a fully connected undirected graph
 - Node \leftrightarrow attribute variable
 - Weight each edge (A_j, A_l) by the empirical MI $\text{MI}(A_j, A_l | C)$
 - Empirical MI = replace chances with relative frequencies
 - Compute the **maximum weight spanning tree**
 - Arbitrarily choose root node
- The TAN T constructed this way maximizes the likelihood $p(\mathbf{d}|T)$

A1 → A2
A3 → A4
A3 → A2
A3 → A1
A3 → A4 → A2
A3 → A1 → A2

14

Learning TAN probabilities

- Focus on mass function $p(A_j | c, a_k, \mathbf{d})$
- Consider the sub-sample related to (c, a_k)
- Apply usual Bayesian learning to the sub-sample

$$p(a_j | c, a_k, \mathbf{d}) = \frac{n_{a_j, c, a_k} + s t_{a_j, c, a_k}}{n_{c, a_k} + s}$$

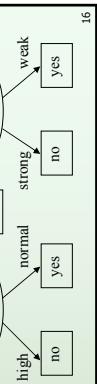
- e.g., $s = 1$ and $t_{a_j, c, a_k} = 1 / |A_j|$ for all a_j
(similarly for the class variable)

The ID3 classification tree

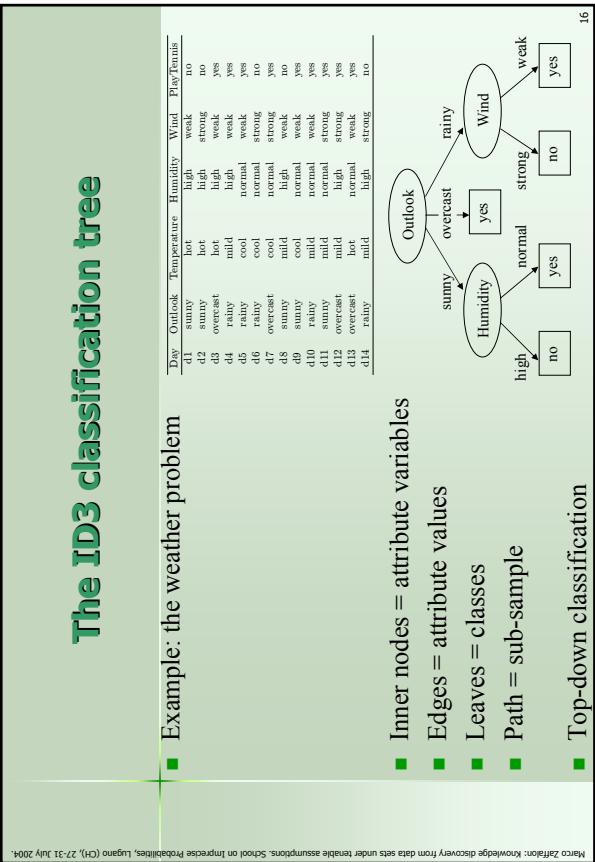
- Example: the weather problem

Day	Outlook	Temperature	Humidity	Wind	PlayTennis
d1	sunny	hot	high	weak	no
d2	sunny	hot	high	strong	yes
d3	overcast	hot	high	weak	yes
d4	rainy	cool	normal	weak	yes
d5	rainy	cool	normal	strong	no
d6	rainy	cool	normal	weak	yes
d7	overcast	normal	high	weak	yes
d8	sunny	normal	high	weak	yes
d9	sunny	normal	high	strong	no
d10	rainy	normal	high	weak	yes
d11	sunny	normal	normal	weak	yes
d12	overcast	normal	normal	strong	yes
d13	overcast	normal	high	weak	yes
d14	rainy	normal	high	strong	no

- Inner nodes = attribute variables
- Edges = attribute values
- Leaves = classes
- Path = sub-sample
- Top-down classification



15



16

ID3 learning

- Compute the empirical MI between class and attribute variables
- Root node = attribute variable with largest MI
- Branch on the variable's attribute values
- Recursion on sub-samples
- Stop criteria
 - No more attribute variables
 - All observations in the same class
 - ...others (avoid overfitting)
- Leaf: class with maximum (empirical) probability

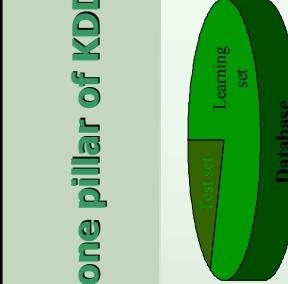


17

Further considerations about KDD

18

Empirical evaluation: one pillar of KDD



- Empirical evaluation of classifiers
 - Split the data into learning and test sets
 - Learn a classifier from the learning set
 - Test it on the test set (hide classes)
- *Prediction accuracy*
 - Relative number of correct predictions
 - e.g., 99% means that 99% of times the predicted and the actual class coincide
- Empirically driven field to some extent
 - Classifier “works” in practice \Rightarrow OK
- Repeated cross validation
 - Randomly partition the data d in the subsets d_1, \dots, d_{10} of approx. equal size
 - For $i = 1 \dots 10$:
 - Learn the classifier from $d \setminus d_i$ and test it on d_i
 - Average the prediction accuracies obtained over the 10 trials
- Cross validation
 - A more sophisticated way to exploit data for testing
- *Tenfold cross validation*

19

Cross validation

20

Serious problems of KDD or ... The necessity of imprecise probability

- The dream of KDD
 - Learning good models from data with minimal, or no, human intervention
- Ignorance matters
 - The learning process is started in conditions of ignorance about the domain
 - This is called *prior ignorance*
 - We may not know some of the values in the data set
 - This is a problem of *missing data*
 - In Bayesian terms
 - Prior ignorance: there is no domain knowledge to choose a prior reliably
 - Missing data are a form of partial ignorance about the likelihood

21

The Bayesian way

- Prior ignorance
 - Model ignorance by so-called *non-informative priors*
 - e.g., Bayes-Laplace's (*uniform*) prior, *Perks' prior*, ...
 - Controversial, long-debated, problem
 - Credible conclusions?
 - e.g., inferred probabilities are precise for any sample size, even equal to zero
 - Problem severe especially for *small* data sets
- Missing data
 - Assume *MAR*, data *missing at random*
 - The mechanism that turns complete into incomplete values is not systematic
 - Missing data can be discarded
 - Problem 1: MAR is often strong and unrealistic
 - Problem 2: MAR cannot be tested statistically
 - Serious problem for any size of data set

22

An imprecise probability approach to prior ignorance

- ## Credibility of conclusions
- Does empirical evaluation help?
 - Prior ignorance
 - Empirical evaluations are unreliable when the sample size is small
 - Missing data
 - Boolean attribute variables, A_1, A_2 . Class variable represents the XOR function, i.e. $C = 1$ iff exactly one of the attribute variables is equal to 1.
 - In the available sample, A_2 is missing iff $A_2 = 0$. The prediction accuracy of the classifier is 100% on the pattern $(A_1=1, A_2=*)$, i.e., the classifier learns to predict 1 all the times.
 - When put to work in practice, the classifier only faces cases s.t. A_2 is missing iff $A_2 = 1$. The accuracy of the classifier drops to 0%. (This is a case of selectively missing data.)
 - \Rightarrow Empirical evaluation can be unreliable with missing data!
 - The law of *decreasing credibility* (C. F. Manski):
 - “the credibility of inference decreases with the strength of the assumptions maintained”
 - Assumptions should be weak enough to be tenable

23

The imprecise Dirichlet model (IDM)

- An imprecise probability method for multinomial sampling
 - Generalization of Bayesian approach
 - Model prior ignorance by the set of all the Dirichlet(s, θ) priors with s fixed
 - s is a *degree of caution*
 - From this, a set of posterior densities
 - Lower and upper probabilities

$$p(\theta) \propto \prod_{d \in \mathcal{D}} \theta_d^{s+d-1}$$

$$\sum_{d \in \mathcal{D}} t_d = 1$$

$$0 < t_d < 1 \quad \forall d \in \mathcal{D}$$

$$p(c, x | \mathbf{d}, s, \mathbf{t}) = \frac{n_{c,x} + st_{c,x}}{N+s}$$

$$p(c, x | \mathbf{d}, s) \in \left[\frac{n_{c,x}}{N+s}, \frac{n_{c,x} + s}{N+s} \right]$$

- Imprecision (upper – lower) is $\frac{s}{N+s}$

25

Credal classification

- There is generally no single *optimal* class with imprecise models
 - e.g., $p(c'|x) \in [0.5, 0.8], p(c''|x) \in [0.3, 0.6]$
 - No class *dominates* the other (they are *incomparable*); both are plausible options
 - Opposite approach: rather than looking for the optimal class, the focus is on **discarding the dominated classes**
 - Credal dominance (strict preference)*
 - $c' > c''$ iff $p(c'|d, x, s, t) > p(c''|d, x, s, t)$ for all t in the IDM
 - iff $\inf_t p(c', x | \mathbf{d}, s, t) / p(c'', x | \mathbf{d}, s, t) > 1$
 - Credal classification*
 - Compare all the classes and output the undominated ones
 - The output is a set of classes! The set shrinks with sample size, reliability

26

The naive credal classifier (NCC)

- Recall that $p(c, x | \mathbf{d}, s, t) = \frac{n_c + st_{c,x}}{N+s} \prod_{j=1}^m \frac{n_{c,a_j} + st_{c,a_j}}{n_{a_j} + st_{a_j}}$
- The optimization problem for testing credal dominance is

$$\inf_{\mathbf{t}} \frac{p(c', x | \mathbf{d}, s, \mathbf{t})}{p(c'', x | \mathbf{d}, s, \mathbf{t})} = \inf_{\mathbf{t}} \left(\frac{n_{c'} + st_{c'}}{n_{c''} + st_{c''}} \right)^{m-1} \prod_{j=1}^m \frac{n_{c',a_j} + st_{c',a_j}}{n_{c'',a_j} + st_{c'',a_j}}$$

$$\text{s.t. } \sum_{c \in \mathcal{C}} t_c = 1$$

$$\begin{aligned} 0 &< t_c < 1 \quad \forall c \in \mathcal{C} \\ 0 &< t_{c,a_j} < t_c \quad \forall c \in \mathcal{C}, j \in \{1, \dots, m\} \\ &= \inf_{\mathbf{t}} \left[\frac{(n_{c'} + st_{c'})^{m-1}}{(n_{c''} + st_{c''})^m} \right] \prod_{j=1}^m \frac{n_{c',a_j}}{n_{c'',a_j} + st_{c'',a_j}} \end{aligned}$$

$$\text{s.t. } t_{c'} + t_{c''} = 1$$

- The latter is a convex optimization problem in a single variable!
- Global optimum found rapidly via numerical method, e.g., Newton-Raphson

27

Some credal classifiers

The "grass grub" application

- Agricultural problem
 - Grass grubs are one of the major insect pests of pasture in Canterbury, New Zealand
 - Objective is to qualitatively predict the grass grub quantity based on characteristics of the paddock and on farming practice
- Data set
 - 155 observations
 - 9 attribute variables
 - Damage ranking, dry or irrigated paddock, position of the paddock, ...
 - Class variable
 - Amount of grass grubs per square meter: low (l), average (a), high (h), very high (v)
 - Relative frequency of the majority class (low) is 0.316

29

NBC performs well

- Tenfold cross-validated performances of several classifiers

Classifier	Accuracy (%)
Decision Table	40.00
IB5	45.16
J48	42.58
Naive Bayes	49.03
OneR	45.16
PART	36.77
SMO	40.64

30

- The prediction accuracy tells only part of the story ...

Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

NBC vs NCC

- Tenfold cross validation
- NBC vs NCC
 - In 60% of cases, NCC outputs a single class, with accuracy 52%
 - In the rest, NCC outputs ~2.3 classes on average (out of 4)
 - Actual class is in this set 82% of times
 - Robust way to deal with scarce knowledge
 - NBC
 - A row, a prior
 - N is the accuracy of the NBC
 - N is the accuracy of the NBC, restricted to the observations where NCC produces indeterminate classifications
 - R_s is the accuracy of a uniformly random predictor
 - NBC is guessing at random when NCC suspends judgment

Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

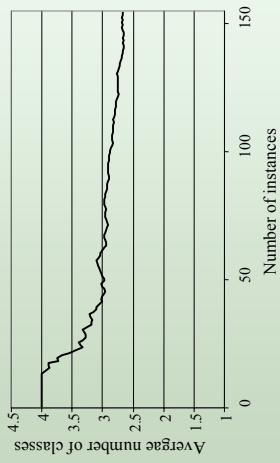
From another angle

- Sequential learning
- Evaluation of NBC performance based on probabilities
 - Loss = logarithmic score
 - $= -\log p(c|\mathbf{d}, \mathbf{x})$
 - Low probability, high loss
 - NBC's total loss = 38.09
 - Total loss of uniformly random predictor = 30
 - Probability = $1/4$ every time

31

32

NCC's sequential learning



33

Concluding remarks on the NCC

- The NCC is reliable and fast
 - Computational complexity
 - $O(N)$ for learning
 - $O(m|C|)$ for classification
 - As fast as NBC
- Open issues
 - e.g., NCC sometimes too cautious
 - Not clear yet why
 - NCC sensitive to irrelevant attribute variables
 - Feature selection probably a necessary step

34

Tree-augmented naive credal classifier (TANC)

- Straightforward extension of TAN to credal classification
- Structure learning as in the precise case
- IDM-based learning of probabilities
- Replace $p(a_j|c, a_k, \mathbf{d}, s, t) = \frac{n_{a_j,c,a_k} + s t_{a_j,c,a_k}}{n_{c,a_k} + s}$ with

$$p(a_j|c, a_k, \mathbf{d}, s) \in \left[\frac{n_{a_j,c,a_k}}{n_{c,a_k} + s}, \frac{n_{a_j,c,a_k} + s}{n_{c,a_k} + s} \right]$$

that is, $p(a_j|\pi_{A_j}) \in [p(a_j|\pi_{A_j}), \bar{p}(a_j|\pi_{A_j})]$

- This way makes different credal sets be *separately specified*
 - (Conditional) Credal sets in the same node and in different nodes

35

TANC classification

- $\underline{p}(c, a_1, \dots, a_m) = \underline{p}(c) \prod_{j=1}^m \underline{p}(a_j|\pi_{A_j})$
- Credal dominance test becomes (complete instance)

$$\frac{\underline{p}(c', a_1, \dots, a_m)}{\bar{p}(c'', a_1, \dots, a_m)} = \frac{\underline{p}(c')}{\bar{p}(c'')} \prod_{j=1}^m \frac{\underline{p}(a_j|\pi'_{A_j})}{\bar{p}(a_j|\pi''_{A_j})} > 1$$
 - Linear time
- More involved with incomplete (ignorable) instances
 - $\underline{p}(c, a_5)?$
 - A_1, \dots, A_4 must be “marginalized out”
 - The expression to minimize involves a sum
 - Exact minimization by propagating intervals over the tree
 - Linear time again

36

Some numbers

Empirical evaluation of TANCs

Data set name (and size)	M%	C ₁ %	C _s %	T %	R _s %	S %	P _s
Iris (150)	33.3	97.0	100.0	92.0	52.9	30.4	11.3
Splice (3170)	51.9	97.7	97.8	94.7	70.4	47.1	11.1
Vehicle (946)	25.4	85.1	85.8	73.6	63.2	36.8	52.4

- M = relative frequency of the majority class
- C₁ = accuracy of the TANC when precise
- C_s = set-based accuracy of the TANC
- T = accuracy of the TAN
- T_s = accuracy of the TAN in the area of imprecision
- R_s = accuracy of the uniformly random predictor in the area of imprecision
- S = size of the area of imprecision
- P_s = average number of classes produced by the TANC in the area of imprecision

- Good, but sometimes too cautious (?)

37

Concluding remarks on the TANC

- Encouraging results
- Sometimes too cautious (?)
- A single IDM?
 - More difficult solving of the optimization problems
- Irrelevant feature variables
 - Problem probably more severe here than with the NCC \Rightarrow feature selection
- Extension to non-ignorable missing data
- Extension of structure learning to imprecise probabilities

38

The problem of incomplete data: an introductory example

- C and A are Boolean random variables
 - C = 1 is the presence of a disease
 - A = 1 is the positive result of a medical test
- Let us do diagnosis
- Good point: you know that
 - $p(C = 0, A = 0) = 0.99$
 - $p(C = 1, A = 1) = 0.01$
- Whence $p(C = 0 \mid A = a)$ allows you to make a sure diagnosis
- Bad point: the test result can be missing
 - This is an incomplete, or *set-valued*, observation {0,1} for A

What is $p(C = 0 \mid A \text{ is missing})?$

Example ctd

- Kolmogorov's *definition* of conditional probability seems to says
 - $p(C = 0 | A \in \{0,1\}) = p(C = 0) = 0.99$
 - i.e., with high probability the patient is healthy
- Is this right?
- In general, it is not
- Why?

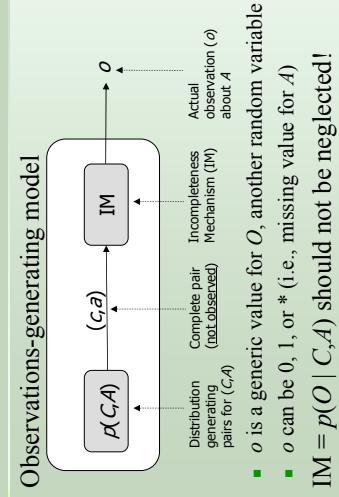
41

Why?

- Because A can be **selectively** reported
 - e.g., the medical test machine is broken; if produces an output \Leftrightarrow the test is negative ($A = 0$)
 - In this case $p(C = 0 | A \text{ is missing}) = p(C = 0 | A = 1) = 0$
 - The patient is definitely ill!
- Compare this with the naive application of Kolmogorov's updating

42

Modeling it the right way



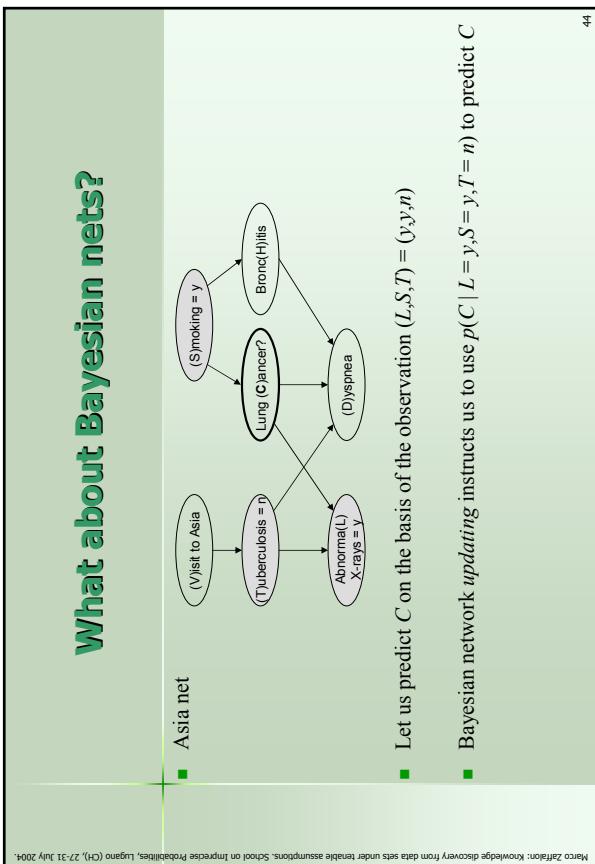
The correct *overall* model we need is $p(C,A)p(O | C,A)$

43

What about Bayesian nets?

- Asia net
- Let us predict C on the basis of the observation $(L,S,T) = (y,y,n)$
- Bayesian network *updating* instructs us to use $p(C | L = y, S = y, T = n)$ to predict C

44



45

Asia ctd

- Should we really use $p(C | L = y, S = y, T = n)$ to predict $C?$

(V, H, D) is missing



$(L, S, T, V, H, D) = (y, y, n, *, *, *)$ is an incomplete observation

- $p(C | L = y, S = y, T = n)$ is just the “naive” updating
- By using the naive updating, we are neglecting the IM!



Wrong inference in general

45

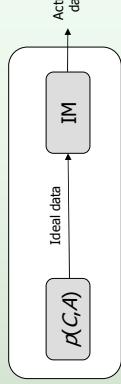
What's the problem with the IM?

- Actually, it can be neglected, if it does not act systematically
 - i.e., if CAR/MAR holds: $p(o|c, a) = a$ for each (c, a)
- Mainstream assumption in literature
- CAR/MAR is very strong and cannot be tested statistically
- Why not modeling the IM explicitly?
 - Often very difficult/costly
 - Partly, because we are not really dealing with “mechanisms” but with humans!
- In many real cases we are left with ignorance about the IM
 - No way but making ignorance become part of our models

46

Statistical treatment of incomplete data

- Pervasive problem in statistical practice, important theoretical issue
- Subtleties
- Fundamental distinction



- Ideal data
 - Produced by a certain process
 - Directly observable
- Actual data
 - Sci-based view of ideal data, produced by the incompleteness mechanism
 - e.g. Space of possibilities = {1, 2, 3}; ideal data = 2; actual data = {1, 2, 3}; missing data = {1, 2, 3}
 - Observable

- Learning and classification with incomplete data
 - Learning set and/or observation to classify incomplete
 - Focus on very general framework
 - e.g., no i.i.d. process

47

Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

48

Learning – ideal variables

- Ideal learning data are regarded as an instance of the random matrix

$$\begin{bmatrix} C_1 & A_{11} & \cdots & A_{1j} & \cdots & A_{1k} \\ \vdots & \vdots & & \vdots & & \vdots \\ C_i & A_{i1} & \cdots & A_{ij} & \cdots & A_{ik} \\ \vdots & \vdots & & \vdots & & \vdots \\ C_N & A_{N1} & \cdots & A_{Nj} & \cdots & A_{Nk} \end{bmatrix} = \begin{bmatrix} C_1 & X_1 \\ \vdots & \vdots \\ C_i & X_i \\ \vdots & \vdots \\ C_N & X_N \end{bmatrix} = \begin{bmatrix} D_1 \\ \vdots \\ D_i \\ \vdots \\ D_N \end{bmatrix} = \mathbf{D}$$

- C_i 's are *class variables*, with values in \mathcal{C}
- A_{ij} 's are *attribute variables*, with values in \mathcal{A}_j for each (i, j)
- Each row D_i in the matrix is called *unit*, with values in \mathcal{A}_m
- $\mathcal{D} = \mathcal{C} \times \mathcal{X} = \mathcal{C} \times \mathcal{A}_1 \times \cdots \times \mathcal{A}_m$

49

Learning – actual variables

- Consider one more random variable
- \mathbf{o} is the *actual observation* of \mathbf{D}
- \mathbf{d} takes values \mathbf{d} from \mathcal{D}^N
- \mathbf{o} takes values \mathbf{o} from $\wp(\mathcal{D}^N)$
- Rather than observing \mathbf{d} , you observe a set \mathbf{o} that contains it
 - This is called *coarsening*, i.e., looking at \mathbf{d} with different levels of detail
 - Note that \mathbf{o} is simply a symbol, not a set, when regarded as the value of \mathbf{o}

50

Learning – assumptions

- Overall model = joint density $p(T, \mathbf{D}, \mathbf{O})$
 - T is a random parameter with values $\theta \in \Theta$
- Factorization
 - $p(T, \mathbf{D}, \mathbf{O}) = p(T)p(\mathbf{D}|T)p(\mathbf{O}|\mathbf{D})$
 - $p(T)$ is an imprecise prior for T , i.e., it belongs to a certain non-empty set
 - The IM depends only on \mathbf{D}
 - Accuracy (of mechanism)
 - $p(\mathbf{d}|\theta) = 0$ if $\mathbf{d} \neq \mathbf{o}$
 - Connection between complete and incomplete observations
 - Positivity
 - Ideal data: $p(\mathbf{D}) > 0$
 - Actual observation: $p(\mathbf{o}) > 0$
 - There exists \mathbf{d} s.t. $p(\mathbf{o}|\mathbf{d}) > 0$

51

Learning – problem

- Formulation of the learning problem
 - Using observed data \mathbf{o} to update beliefs about a function $f : \Theta \rightarrow \mathbb{R}$
 - e.g., θ real number, $f(\theta) = \theta$
 - In the precise framework, one would compute $E(f|\mathbf{o})$
- $E(f|\mathbf{o}) = \frac{\int f(\theta)p(\theta, \mathbf{o})d\theta}{p(\mathbf{o})}$
 $= \frac{\int f(\theta)\sum_{\mathbf{d} \in \mathcal{O}}p(\theta)p(\mathbf{d}|\theta)p(\mathbf{o}|\mathbf{d})d\theta}{\sum_{\mathbf{d} \in \mathcal{O}}p(\mathbf{o}|\mathbf{d})p(\mathbf{d})}$
- Using Factorization, Accuracy and Positivity
- What about $p(\mathbf{o}|\mathbf{d})$ and $p(\theta)$?
 - Imprecise knowledge

52

Learning – express ignorance about the IM

- Focus on $p(\boldsymbol{o}|\boldsymbol{d})$, $\boldsymbol{d} \in \boldsymbol{o}$
 - Call $p(\boldsymbol{o}|\boldsymbol{D})$ the vector with elements $p(\boldsymbol{o}|\boldsymbol{d})$, $\boldsymbol{d} \in \boldsymbol{o}$
 - We can only constrain the set of possible vectors $p(\boldsymbol{o}|\boldsymbol{D})$
- $\sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) > 0$
- $0 \leq p(\boldsymbol{o}|\boldsymbol{d}) \leq 1$, $\boldsymbol{d} \in \boldsymbol{o}$
- The restriction on the sum is due to Accuracy
 - The inequality following the sum is due to Positivity
 - The inequalities define an open linear set called $\mathcal{P}(\boldsymbol{o}|\boldsymbol{D})$
 - Call $\mathcal{P}_e(\boldsymbol{o}|\boldsymbol{D})$ the approximating closed set
- $\sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) \geq \varepsilon$
- $0 \leq p(\boldsymbol{o}|\boldsymbol{d}) \leq 1$, $\boldsymbol{d} \in \boldsymbol{o}$

53

Learning – goal

- Goal: $\underline{E}(f|\boldsymbol{o}) = \inf_{p(\boldsymbol{T}) \in \mathcal{P}(\boldsymbol{T}), p(\boldsymbol{o}|\boldsymbol{D}) \in \mathcal{P}(\boldsymbol{o}|\boldsymbol{D})} \inf_{E(f|\boldsymbol{o})} E(f|\boldsymbol{o})$
- Focus on $\min_{p(\boldsymbol{o}|\boldsymbol{D}) \in \mathcal{P}_e(\boldsymbol{o}|\boldsymbol{D})} E(f|\boldsymbol{o})$
 - i.e., $\min_{p(\boldsymbol{o}|\boldsymbol{D}) \in \mathcal{P}_e(\boldsymbol{o}|\boldsymbol{D})} \frac{\sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) \int f(\boldsymbol{\theta}) p(\boldsymbol{\theta}) n(d|\boldsymbol{\theta}) d\boldsymbol{\theta}}{\sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) p(\boldsymbol{d})}$
 - Objective function is ratio of linear functions
- Fractional programming theorem
 - Consider $\min_{x \in S} \frac{q(x)}{r(x)}$ where $S = \text{compact subset of } \mathbb{R}^v$
 - q, r continuous and r positive on S
 - Define $h : \mathbb{R} \rightarrow \mathbb{R}$, $h(\mu) = \min_{x \in S} [q(x) - \mu r(x)]$
 - Then $\mu^* = \arg\min_{x \in S} \frac{q(x)}{r(x)}$ $\iff h(\mu^*) = 0$
 - h has a single zero

54

Learning – solution: conservative learning rule

$$\begin{aligned}
 h(\mu) &= \min_{p(\boldsymbol{o}|\boldsymbol{D}) \in \mathcal{P}_e(\boldsymbol{o}|\boldsymbol{D})} \left[\sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) \int f(\boldsymbol{\theta}) p(\boldsymbol{\theta}) p(\boldsymbol{d}|\boldsymbol{\theta}) d\boldsymbol{\theta} - \mu \sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) p(\boldsymbol{d}) \right] \\
 &= \min_{p(\boldsymbol{o}|\boldsymbol{D}) \in \mathcal{P}_e(\boldsymbol{o}|\boldsymbol{D})} \sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) p(\boldsymbol{d}) \left[\int f(\boldsymbol{\theta}) p(\boldsymbol{\theta}|\boldsymbol{d}) d\boldsymbol{\theta} - \mu \right] \\
 &= \min_{p(\boldsymbol{o}|\boldsymbol{D}) \in \mathcal{P}_e(\boldsymbol{o}|\boldsymbol{D})} \sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) p(\boldsymbol{d}) [E(f|\boldsymbol{d}) - \mu] \\
 \mu^* &= \min_{\boldsymbol{d} \in \boldsymbol{o}} E(f|\boldsymbol{d}) \implies h(\mu^*) = 0 \\
 \min_{p(\boldsymbol{o}|\boldsymbol{D}) \in \mathcal{P}_e(\boldsymbol{o}|\boldsymbol{D})} E(f|\boldsymbol{o}) &= \min_{\boldsymbol{d} \in \boldsymbol{o}} E(f|\boldsymbol{d})
 \end{aligned}$$



$$\underline{E}(f|\boldsymbol{o}) = \inf_{p(\boldsymbol{T}) \in \mathcal{P}(\boldsymbol{T})} \min_{\boldsymbol{d} \in \boldsymbol{o}} E(f|\boldsymbol{d})$$

Classification – variables

- Learning variable \boldsymbol{D} as before
- The ideal observation to classify is regarded as a (partial) instance of the further unit $(C, A_1, \dots, A_j, \dots, A_N) = (C, X) = D$ with values in \mathcal{D}
- Summary of ideal variables

$$\begin{bmatrix} C \\ A_1 \\ A_2 \\ \vdots \\ A_k \\ \vdots \\ A_N \end{bmatrix} \boldsymbol{D} + \begin{bmatrix} C \\ A_1 \\ A_2 \\ \vdots \\ A_k \\ \vdots \\ A_N \end{bmatrix} = \begin{bmatrix} C \\ A_1 \\ A_2 \\ \vdots \\ A_k \\ \vdots \\ A_N \\ A_N \\ A_N \\ \vdots \\ A_N \\ A_N \\ \vdots \\ A_N \end{bmatrix}$$

- Actual variables: $\boldsymbol{o}, \boldsymbol{o}^+, \boldsymbol{o}^-$
- Taking values in the related power sets

55

56

Classification – assumptions

- Factorization, Accuracy, Positivity, as before
- One more: Independence
 - $p(O^+|D^+) = p(O^+|D^-)$
 - The mechanism does not depend on what we want to predict, i.e., C
 - Without it, we could never exclude that $p(c|o^+) = 0$ or $p(c|o^-) = 1$
 - Independence avoids problem of vacuous conclusions
- On the meaning of Independence
 - Equivalent to $p(C|O^+, D^-) = p(C|D^-)$
 - Independence characterizes problems of incomplete data, in the sense that
 - Once you know D^- , there is no more a problem of incomplete (or missing) data
 - Consider the opposite case: you have not included some factor
 - Independence not very restrictive in practice
 - If D^- represents the set of all "factors" that you deem important to predict C , Independence follows automatically

57

Classification – goal and solution: conservative updating rule

- Focus on $E_C(g|o^+) = \sum_{c \in C} g(c)p(c|o^+)$
- Generic g
 - e.g., with credal dominance: $g(c) = \begin{cases} 1 & c = c' \\ -1 & c = c'' \\ 0 & \text{otherwise} \end{cases}$
- Goal: $\underline{E}_C(g|o^+)$
- Theorem: $\underline{E}_C(g|o^+) = \inf_{p(T) \in \mathcal{P}(T)} \min_{d^- \in o^-} E_C(g|d^-)$
- Incomplete data, via the *conservative updating rule*,
 - naturally produce credal classifiers!

58

Obtaining stronger conclusions: the mixed rules

- CAR might hold in some cases
 - This should not be neglected in order to strengthen conclusions
- Some attribute variables are CAR, the others non-ignorable
- Mixed rules

- Learning: $\underline{E}(f|o, \hat{o}) = \inf_{p(T) \in \mathcal{P}(T)} \min_{d \in o} E(f|d, \hat{o} \in \hat{o})$
- Classification: $\underline{E}_C(g|o^+, \hat{o}^+) = \inf_{p(T) \in \mathcal{P}(T)} \min_{d^- \in o^-} E_C(g|d^-, \hat{d}^- \in \hat{o}^-)$

The conservative rules in practice

Conservative learning in the unstructured case

Unstructured case

- Table with missing data
- Focus on the empirical mass function $p(X)$
- and on derived descriptive indexes
- Conservative learning:
 - Consider all the complete tables (completions)
 - consistent with the incomplete one
- Problem of complexity
 - n missing values $\Rightarrow 3^n$ completions
 - 3^5 in the example

62

Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

Obs.	O_{i1}	O_{i2}
o_1	1	1
o_2	1	*
o_3	1	*
o_4	*	2
o_5	3	2
o_6	*	*

$$(A_1, A_2) = X \text{ in } \{1, 2, 3\}^2$$

$$(O_{A_1}, O_{A_2}) = O \text{ in } \{1, 2, 3, *\}^2$$

Alternative representation

- Data set \Rightarrow network of flow

Obs.	O_{i1}	O_{i2}
o_1	1	1
o_2	1	*
o_3	1	*
o_4	*	2
o_5	3	2
o_6	*	*

Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

Flow network = polytope

- The possible flows make up a polytope
- Let $R(o_i) =$ set of nodes to which o_i is mapped
- $f \in \Psi$ where Ψ is the defined by

$$\sum_{x \in R(o_i)} f(o_i, x) = 1 \quad i = 1, \dots, N$$

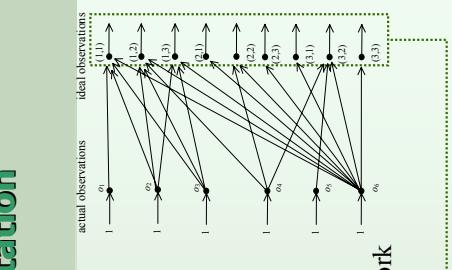
$$\sum_{j: x \in R(o_j)} f(o_j, x) = f(x, \cdot) \quad \forall x \in \mathcal{X}$$

$$f \geq 0$$

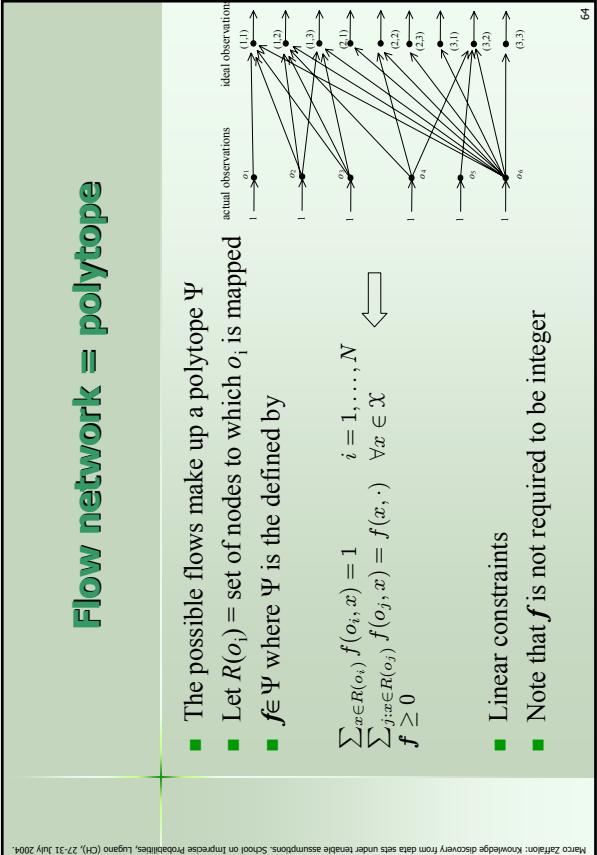
- Linear constraints
- Note that f is not required to be integer

63

Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.



- Let the vector f be a given flow in the network
- With elements $f(a,b)$: flow on the arc $a \rightarrow b$
- Denote by f_X the sub-vector for these arcs
- $p(X) = (1/N)f_X$



Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

Alternative representation

- ideal observations
- actual observations
- The possible flows make up a polytope
- Let $R(o_i) =$ set of nodes to which o_i is mapped
- $f \in \Psi$ where Ψ is the defined by

$$\sum_{x \in R(o_i)} f(o_i, x) = 1 \quad i = 1, \dots, N$$

$$\sum_{j: x \in R(o_j)} f(o_j, x) = f(x, \cdot) \quad \forall x \in \mathcal{X}$$

$$f \geq 0$$

$$\sum_{x \in R(o_i)} f(o_i, x) = 1 \quad i = 1, \dots, N$$

$$\sum_{j: x \in R(o_j)} f(o_j, x) = f(x, \cdot) \quad \forall x \in \mathcal{X}$$

$$f \geq 0$$

- Linear constraints
- Note that f is not required to be integer

64

Marco Zaffalon: Knowledge discovery from data sets under testable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

Properties of the flow model

- $\tilde{\mathcal{P}}$ = finite set of joint mass functions from all the completions
- $\tilde{\mathcal{P}}_f$ = finite set of joint mass functions from all the integer flows
- \mathcal{P} and \mathcal{P}_f are their respective convex hulls

Lemma 1: $\tilde{\mathcal{P}} = \tilde{\mathcal{P}}_f$

Lemma 2: The extreme points of Ψ are integer flows

Theorem: $\mathcal{P} = \mathcal{P}_f = \left\{ \frac{1}{N} \mathbf{f}_X \mid \mathbf{f} \in \Psi \right\}$

- The constraints for Ψ provide us with an implicit description of \mathcal{P}
- No need to consider the extreme distributions explicitly

65

Linear-programming computations

- $\mathcal{X}', \mathcal{X}'' \subseteq \mathcal{X}$
- Optimizing linear functions
 - e.g., $\underline{p}(\mathcal{X}'), \bar{p}(\mathcal{X}'')$
 - Exact computation, polynomial time
- Optimizing special types of fractional linear functions
 - e.g., $\underline{p}(\mathcal{X}'|\mathcal{X}''), \bar{p}(\mathcal{X}'|\mathcal{X}'')$
 - Exact computation, polynomial time
- Optimizing general fractional linear functions
 - e.g., $\underline{E}(\mathcal{X}'|\mathcal{X}''), \bar{E}(\mathcal{X}'|\mathcal{X}'')$
 - Approximate computation (any precision), polynomial time

66

Further properties

- \underline{p} is a belief function
 - Special algorithms on the network lead to linear-time computations
- Modeling more general patterns of missing data
 - Partially missing type 1
 - Not all values are possible replacements
 - Partially missing type 2
 - Intra-observation dependencies
- Easy extension to an inferential approach
 - imprecise Dirichlet model + treatment of missing data
 \Downarrow
 - prior ignorance + partial ignorance about the likelihood
- e.g., $[\underline{p}(\mathcal{X}'|\mathbf{o}, s), \bar{p}(\mathcal{X}'|\mathbf{o}, s)] = \left[\frac{\underline{n}(\mathcal{X}')}{N+s}, \frac{\bar{n}(\mathcal{X}')+s}{N+s} \right]$

67

1

2

3

The naively structured case

4

5

6

68

Inference of the NCC from incomplete data

- Test of credal dominance with an incomplete learning sample

$$\min_{\mathbf{d} \in \Theta} \frac{p(c', x | \mathbf{d}, s, t)}{p(c'', x | \mathbf{d}, s, t)} = \inf_{t \in \Theta} \left[\frac{(n_{c'} + st_{c'})}{(n_{c'} + st_{c''})} \right]^{m-1} \prod_{j=1}^m \frac{n_{c'} a_j}{\bar{n}_{c''} a_j + st_{c''}}$$

s.t. $t_{c'} + t_{c''} = 1$
 $t_{c'}, t_{c''} > 0$

- Same complexity as with complete data

The dementia application

- Real problem of diagnosing dementias

- Application of the NCC
 - Inference from incomplete data
 - Conservative learning

Aim

- To show that developed methods are useful
- To compare them with more traditional methods
- To show that developed methods can be used in practice right now

70

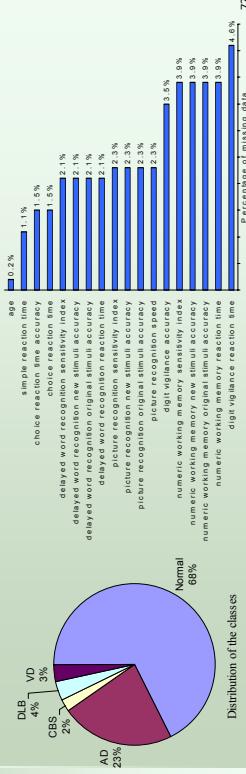
The Cognitive Drug Research (CDR) system

- Assessment of cognitive functions
- Most widely used system in clinical research
- Computer-based automated system
- CDR tasks and functions assessed
 - Attention, concentration and vigilance
 - Simple reaction time, choice reaction time, digit vigilance
 - Working memory and executive function
 - Episodic secondary memory
 - Picture recognition

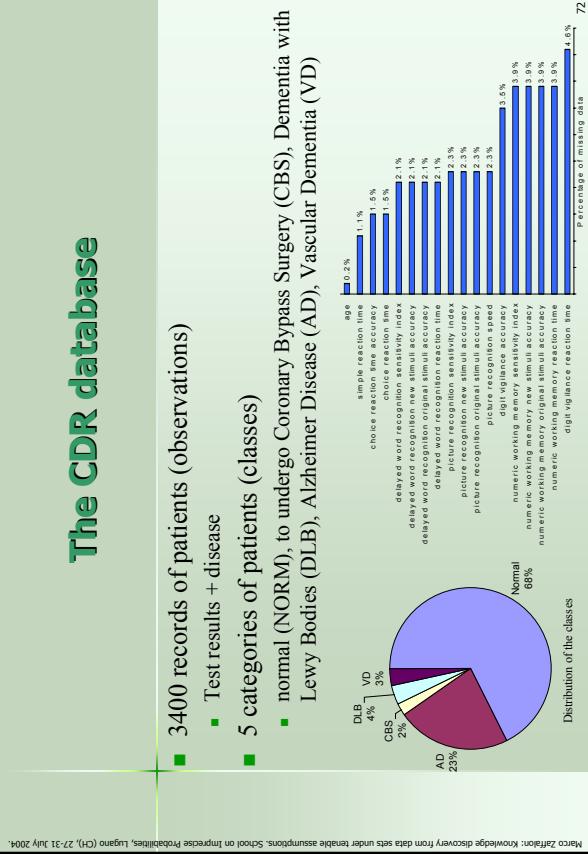
Marco Zaffalon: Knowledge discovery from data sets under unreliable assumptions, School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

The CDR database

- 3400 records of patients (observations)
 - Test results + disease
- 5 categories of patients (classes)
 - normal (NORM), to undergo Coronary Bypass Surgery (CBS), Dementia with Lewy Bodies (DLB), Alzheimer Disease (AD), Vascular Dementia (VD)



71



72

Differentiating dementias

- Placing a demented patient in the right class
- Training and test data
 - ~50% - 50%
- Results
 - | C ₁ % | C _S % | N% | N _S % | R _S % | S% | P _S % |
|------------------|------------------|-------|------------------|------------------|-------|------------------|
| 94.05 | 98.42 | 89.76 | 75.59 | 45.6 | 23.22 | 2.3/4 |
- C₁ = accuracy of the NCC when precise
- C_s = set-based accuracy of the NCC
- N = accuracy of the NCC
- N_s = accuracy of the NCC in the area of imprecision
- R_s = accuracy of the uniformly random predictor in the area of imprecision
- S = size of the area of imprecision
- P_s = average number of classes produced by the NCC in the area of imprecision

73

Other cases

74

Robust classification with TANCS

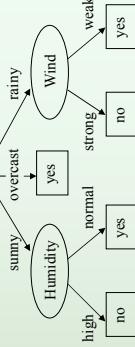
- TANCS can be extended to classifying non-ignorable missing data
 - Case of ignorable missing data already there
 - New extension = conservative updating with a complete learning set
 - As the Markov blanket of the class node is singly connected (a tree)
 - Work in progress
 - Extension to conservative learning
 - This is more difficult
 - Missing data create non-separately specified credal sets

Marcos Zaffalon: Knowledge discovery from data sets under testable assumptions. School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

The case of ID3

75

- ID3 can be extended to classifying non-ignorable missing data
 - New extension = conservative updating with a complete learning set
 - How: follow all the paths downward
- Learning is more difficult
 - Many trees



Marcos Zaffalon: Knowledge discovery from data sets under testable assumptions. School on Imprecise Probabilities, Lugano (CH), 27-31 July 2004.

76

Concluding remarks

- Credible conclusions need tenable assumptions
 - Even if empirical validations are possible!
- Imprecise probabilities permit working with weak assumptions
 - Prior ignorance, incomplete data
- Methods exist that already work in practice
 - Useful
 - Efficient
- More work is (always) needed
 - Creating more methods
 - Developing efficient implementations

77

Annotated bibliography

- Introduction to pattern classification
 - R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern classification*. Wiley, 2001. 2nd edition.
- Bayesian data analysis
 - A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman, 1995.
- First definition of the NBC
 - R. O. Duda and P. E. Hart. *Pattern classification and scene analysis*. Wiley, New York, 1973.
- TAN classifier
 - N. Friedman, D. Geiger, and M. Goldszmidt. Bayesian networks classifiers. *Machine Learning*, 29(2/3):131–163, 1997.

78

1

Annotated bibliography

- Definition of credal classification
 - M. Zaffalon. The naive credal classifier. *Journal of Statistical Planning and Inference*, 105(1):5–21, 2002.
- Inference of the NCC from (possibly incomplete non-ignorable) data
 - M. Zaffalon. Statistical inference of the naive credal classifier. In G. de Cooman, T. L. Fine, and T. Seidenfeld, editors, *ISIPTA '01: Proceedings of the Second International Symposium on Imprecise Probabilities and Their Applications*, pages 384–393, The Netherlands, 2001. Shaker.
- Grass grub application
 - M. Zaffalon. Credible classification for environmental problems. *Environmental modelling and software*. To appear.
- TANC classifier
 - M. Zaffalon and E. Fagioli. Tree-based credal networks for classification. *Reliable Computation*, 9(6):487–509, 2003.

80

2

1

Annotated bibliography

- Classification trees
 - J. R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, 1993.
- The classical book on (ignorable) missing data
 - R. J. A. Little and D. B. Rubin. *Statistical Analysis with Missing Data*. Wiley, New York, 1987.
- Imprecise probability
 - P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, New York, 1991.
 - The imprecise Dirichlet model
 - P. Walley. Inferences from multinomial data: learning about a bag of marbles. *J. R. Statist. Soc. B*, 58(1):3–57, 1996.

79

3

Annotated bibliography

- Wide view of incomplete non-ignorable data in identification prob.
 - C. F. Manski. *Partial Identification of Probability Distributions*. Springer-Verlag, New York, 2003.
- Non-ignorable missing data in artificial intelligence
 - M. Ramoni and P. Sebastiani. Robust learning with missing data. *Machine Learning*, 45(2):147–170, 2001.
- Problems of naive updating raised for the first time
 - G. Shafer. Conditional probability. *International Statistical Review*, 53:261–277, 1985.
- Naive updating works iff CAR/MAR; CAR/MAR is strong
 - P. Grünwald and J. Halpern. Updating probabilities. *Journal of Artificial Intelligence Research*, 19:243–278, 2003.

81

Annotated bibliography

- First definition of the conservative updating rule (expert systems)
 - G. de Cooman and M. Zaffalon. Updating beliefs with incomplete observations. *Artificial Intelligence*. Forthcoming.
- Statistical framework for non-ignorable incomplete data
 - M. Zaffalon. Statistical classification with incomplete observations. Technical Report IDSLA-10-04, IDSLA, 2004. In progress.
- Dealing with non-ignorable missing data in the unstructured case
 - M. Zaffalon. Exact credal treatment of missing data. *Journal of Statistical Planning and Inference*, 105(1):105–122, 2002.
- Dementia application
 - M. Zaffalon, K. Wessels, and O. Petruini. Reliable diagnoses of dementia by the naive credal classifier inferred from incomplete cognitive data. *Artificial Intelligence in Medicine*, 29(1–2):61–79, 2003.

82