

Exercises

Exercise 1 – IDM

- Consider the following data set (assume Boolean variables)

#	C	A
1	0	1
2	0	1
3	0	1

- Use the IDM($s=2$) sequentially to
 - compute the lower and upper probabilities of ($C=1|A=1$) for the observations 0, 1, 2, 3
 - classify the same observations
- Repeat the above using the precise Dirichlet model with Bayes-Laplace's uniform prior, and compare the results

2

Exercise 2 – NCC, complete data

- Consider the weather data set

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

- Use the NCC($s=1$) to classify the instance (overcast, cool, high, strong)
- What if Outlook = sunny?

3

Exercise 3 – NCC, incomplete data

- Consider the incomplete weather data set

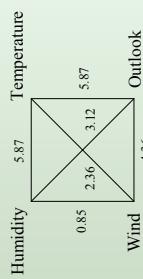
Day	Outlook	Temperature	Humidity	Wind	PlayTennis
d1	sunny	hot	high	high	no
d2	sunny	hot	high	strong	yes
d3	overcast	hot	high	weak	yes
d4	rainy	normal	normal	weak	yes
d5	rainy	normal	normal	weak	yes
d6	rainy	normal	normal	strong	no
d7	overcast	normal	normal	strong	yes
d8	overcast	normal	normal	strong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

- Use the NCC($s=1$) to test whether or not 'no' dominates 'yes' with the instance (sunny, *, high, *)
 - Assume that 'Temperature' is missing at random, while 'Outlook' and 'Wind' are missing in non-i-ignorable way
 - Tip: minimum of function to optimize in $t_{yes} = 1$

4

Exercise 4 – TANC

- Consider the (complete) weather data set
- Use the TANC($s=1$) to classify the instance (sunny, mild, high, strong)
 - Assuming the following values of conditional mutual information
- and taking Humidity as root of the tree
 - and taking Humidity as root of the tree



5

Exercise 5 – mixed rules

- Sketch the theoretical derivation of the mixed rule in the learning setting
 - Stating new assumptions if needed

6

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($N+s, \mathbf{t}^*$): $p(\theta | \mathbf{d}) \propto \prod_{d \in \mathcal{D}} \theta_d^{s t_d + s t_d - 1}$
 - $\mathbf{t}^* = (t_d^*)_{d \in \mathcal{D}}, t_d^* = (n_d + s t_d) / (N + s)$
 - t_d^* = expectation of θ_d w.r.t. the posterior
- The needed probability is then

$$p(c, x | \mathbf{d}) = E[\theta_c | \mathbf{d}] = \frac{n_{c,x} + s t_{c,x}}{N+s}$$

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.

Appendixes

7

8

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, \boldsymbol{t}) priors with s fixed
 - s is a *degree of caution*
 - From this, a set of posterior densities
- Lower and upper probabilities

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

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

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

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

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

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

9

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 \boldsymbol{t} in the IDM
 - iff $\inf_{\boldsymbol{t}} p(c', x | \boldsymbol{d}, s, \boldsymbol{t}) / p(c'', x | \boldsymbol{d}, s, \boldsymbol{t}) > 1$
 - Classification
 - Compare all the classes and output the undominated ones
 - The output is a set of classes! The set shrinks with sample size: reliability
 - Credal classification*

10

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

The naive credal classifier (NCC)

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

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

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

$$0 < t_c < 1 \quad \forall c \in \mathcal{C}$$

$$0 < t_{c,a_j} < t_c \quad \forall c \in \mathcal{C}, j \in \{1, \dots, m\}$$

$$= \inf \left[\frac{\left(\frac{n_{c'}}{n_{c''}} + st_{c'} \right)^{m-1}}{\left(\frac{n_{c'}}{n_{c''}} + st_{c'} \right)^m} \prod_{j=1}^m \frac{n_{c',a_j}}{n_{c'',a_j} + st_{c''}} \right]$$

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

$$t_{c'}, t_{c''} > 0$$
- The latter is a convex optimization problem in a single variable!
- Global* optimum found rapidly via numerical method, e.g., Newton-Raphson

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

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

Relaxing the independence assumption: the TAN model

- Tree augmented naive 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

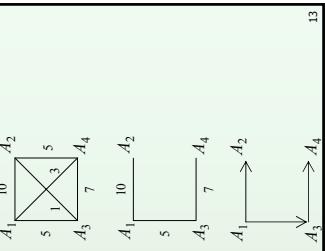
$$p(c, x | \boldsymbol{d}) = p(c | \boldsymbol{d}) \prod_{j=1}^m p(a_j | \pi_{A_j}, \boldsymbol{d})$$

12

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

Learning TAN structure

- Consider the *conditional mutual information* (MI)
$$\text{MI}(A_j, A_k | C) = \sum_{c \in C} \theta_c \sum_{(a_j, a_k) \in A_j \times A_k} \frac{\theta_{a_j, a_k | c} \log \frac{\theta_{a_j, a_k | c}}{\theta_{a_j | c} \theta_{a_k | c}}}{\theta_{a_j | c} \theta_{a_k | c}}$$
- Procedure**
 - Set up a fully connected undirected graph
 - Node \leftrightarrow attribute variable
 - Weight each edge (A_i, A_j) by the *empirical MI* $\text{MI}(A_i, A_j | C)$
 - 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 $p(\mathbf{d}|T)$



13

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, $p(a_j | \pi_{A_j}) \in [\underline{p}(a_j | \pi_{A_j}), \bar{p}(a_j | \pi_{A_j})]$
- This way makes different credal sets be *separately specified*
 - (Conditional) Credal sets in the same node and in different nodes

14

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)
$$\frac{\underline{p}(c', a_1, \dots, a_m)}{\underline{p}(c'', a_1, \dots, a_m)} = \frac{\underline{p}(c')^m}{\underline{p}(c'')} \prod_{j=1}^m \frac{\underline{p}(a_j | \pi'_{A_j})}{\underline{p}(a_j | \pi''_{A_j})} > 1$$
- 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

15

Learning – ideal variables

- Ideal learning data are regarded as an instance of the random matrix
$$\begin{bmatrix} C_1 & A_{11} & \dots & A_{1k} & \dots & A_{1N} \\ \vdots & \vdots & & \vdots & & \vdots \\ C_i & A_{i1} & \dots & A_{ij} & \dots & A_{ik} \\ \vdots & \vdots & & \vdots & & \vdots \\ C_N & A_{N1} & \dots & A_{Nj} & \dots & A_{Nk} \end{bmatrix} = \begin{bmatrix} C_1 & X_1 \\ \vdots & \vdots \\ C_i & X_i \\ \vdots & \vdots \\ C_N & X_N \end{bmatrix} = \begin{bmatrix} D_1 \\ \vdots \\ D_i \\ \vdots \\ D_N \end{bmatrix} = \mathbf{D}$$
- 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 $\mathcal{D} = \mathcal{C} \times \mathcal{X} = \mathcal{C} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_m$

16

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 *censoring*, 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}

17

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)
 - $p(\mathbf{o}|\mathbf{d}) = 0$ 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. $p(\mathbf{o}|\mathbf{d}) > 0$

18

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

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 $E(f|\mathbf{o})$

$$\begin{aligned} E(f|\mathbf{o}) &= \frac{\int f(\theta)p(\theta,\mathbf{o})d\theta}{p(\mathbf{o})} \\ &= \frac{\int f(\theta)\sum_{d \in \mathbf{o}}p(\theta)p(\mathbf{d}|\theta)p(\mathbf{o}|\mathbf{d})d\theta}{\sum_{d \in \mathbf{o}}p(\mathbf{o}|\mathbf{d})p(\mathbf{d})} \end{aligned}$$

- Using Factorization, Accuracy and Positivity
- What about $p(\mathbf{o}|\mathbf{d})$ and $p(\theta)$?
 - Imprecise knowledge

Learning – express ignorance about the IM

- Focus on $p(\mathbf{o}|\mathbf{d})$, $\mathbf{d} \in \mathbf{o}$
 - Call $p(\mathbf{o}|\mathbf{D})$ the vector with elements $p(\mathbf{o}|\mathbf{d})$, $\mathbf{d} \in \mathbf{o}$
- We can only constrain the set of possible vectors $p(\mathbf{o}|\mathbf{D})$
 - $\sum_{d \in \mathbf{o}}p(\mathbf{o}|\mathbf{d}) > 0$
 - $0 \leq p(\mathbf{o}|\mathbf{d}) \leq 1$, $d \in \mathbf{o}$
 - 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 $\mathcal{P}(\mathbf{o}|\mathbf{D})$
 - Call $\mathcal{P}_\varepsilon(\mathbf{o}|\mathbf{D})$ the approximating closed set
 - $\sum_{d \in \mathbf{o}}p(\mathbf{o}|\mathbf{d}) \geq \varepsilon$
 - $0 \leq p(\mathbf{o}|\mathbf{d}) \leq 1$, $d \in \mathbf{o}$

19

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.

Learning – goal

- Goal: $\underline{E}(f|\boldsymbol{o}) = \inf_{p(T) \in \mathcal{P}(T)} \inf_{p(\boldsymbol{o}|D) \in \mathcal{P}(\boldsymbol{o}|D)} E(f|\boldsymbol{o})$
- Focus on $\min_{p(\boldsymbol{o}|D) \in \mathcal{P}_e(\boldsymbol{o}|D)} E(f|\boldsymbol{o})$
 - i.e., $\min_{p(\boldsymbol{o}|D) \in \mathcal{P}_e(\boldsymbol{o}|D)} \frac{\sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) \int f(\theta) p(\theta) p(\boldsymbol{d}|\theta) d\theta}{\sum_{\boldsymbol{d} \in \boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{d}) p(\boldsymbol{d})}$
 - 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 $\mu^* = \operatorname{argmin}_{x \in S} \frac{q(x)}{r(x)} \iff h(\mu^*) = 0$
 - h has a single zero

21

Learning – solution: conservative learning rule

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

22

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

23

Inference of the NCC from incomplete data

- Test of credal dominance with an incomplete learning sample
- $$\min_{\boldsymbol{d} \in \boldsymbol{o}} \frac{p(c', x|\boldsymbol{d}, s, t)}{p(c'', x|\boldsymbol{d}, s, t)} = \inf_{\boldsymbol{d} \in \boldsymbol{o}} \left(\frac{n_{c'} + st_{c'}}{n_{c''} + st_{c''}} \right)^{m-1} \prod_{j=1}^m \frac{n_{c', a_j}}{n_{c'', a_j} + st_{c''}}$$

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

- Same complexity as with complete data

24