Optimal weight constraint elicitation for additive multi-attribute utility models

Gert van Valkenhoef, Tommi Tervonen

EURO 2013, Rome, Italy, July 2013

Setting

Stochastic additive multi-attribute utility models

- Stochastic criteria measurements
- Separate per-criterion evaluation (partial utilities) from relative importance (weights)
- Eliciting preference information (i.e. linear weight constraints)
  - Ordinal SWING (easy, imprecise)
  - SWING (moderately hard, overly precise?)
  - Interval SWING (hard, precise)
  - Pair-wise judgments (moderate, time-precision trade-off)

Setting

Stochastic additive multi-attribute utility models

- Stochastic criteria measurements
- Separate per-criterion evaluation (partial utilities) from relative importance (weights)
- Eliciting preference information (i.e. linear weight constraints)
  - Ordinal SWING (easy, imprecise)
  - SWING (moderately hard, overly precise?)
  - Interval SWING (hard, precise)
  - **Pair-wise judgments** (moderate, time-precision trade-off)

Setting

Stochastic additive multi-attribute utility models

- Stochastic criteria measurements
- Separate per-criterion evaluation (partial utilities) from relative importance (weights)
- Eliciting preference information (i.e. linear weight constraints)
  - Ordinal SWING (easy, imprecise)
  - SWING (moderately hard, overly precise?)
  - Interval SWING (hard, precise)
  - **Pair-wise judgments** (moderate, time-precision trade-off)
- How should we prioritize questions? How much information is enough?

Notation

- Alternatives $I = \{1, \ldots, m\}$
- Criteria $J = \{1, \ldots, n\}$
- Measurements x_j^i with joint density $f_X(x)$
- Partial utility functions $u_j : \mathbb{R} \to [0, 1]$
- Utility

$$u(x^i, w) = \sum_{j \in J} w_j u_j(x_j^i)$$

- Weights w_j are non-negative normalized
- Incomplete information on w: density $f_W(w)$
  - As usual in SMAA
  - Most often uniform over feasible weights

Problem

- Elicit preference information to reduce uncertainty on
  - Best alternative (choice)
  - Ranking of alternatives (ranking)
  - Classification of alternatives (sorting)
- Uncertainty comes from both
  - The criteria measurements' joint distribution $f_X(x)$
  - Incomplete information on weights $f_W(w)$
- In this setting, how do we decide
  - Which question to ask next?
  - When to stop asking questions?

Entropy

Entropy: formal framework for reasoning about information

- Shannon entropy quantifies expected value of information

$$H(X) = E[I(X)] = E[-\log P(X)]$$
$$= -\sum_x p(X = x) \log p(X = x)$$

- $I(X)$ is the *information* or *surprisal*
- Higher entropy means more uncertainty
- Conditional entropy:

$$H(Y|X) = \sum_x p(x) H(Y|X = x)$$

Entropy minimization method

- Define decision entropy
  - For a given problematic (choice, ranking, sorting)
  - Depends on PDF of measurements, weights
- Find the question that minimizes entropy
  - i.e. maximizes expected information gain
  - conditional on the answer
- Stopping criterion
  - entropy is known within a certain band width

Decision entropy: choice problem

The decision uncertainty (entropy) given feasible weights W':

$$H(Y|W') = -\sum_{y} p(y|W') \log p(y|W')$$

Where $y \in I$ is the alternative that obtains first rank:

$$p(y|W') = \int_{w \in W_n} \int_{x \in \mathbb{R}^{m \times n}} f_W(w|w \in W') f_X(x) p(y|x, w) \, \mathrm{d}x\mathrm{d}w$$

$$p(y|x, w) = \begin{cases} 1 & \text{if } \mathrm{rank}(y|x, w) = 1 \\ 0 & \text{otherwise} \end{cases}$$

Question-answer entropy

- Set of possible questions $\mathbf{Q}$
- For each $Q \in \mathbf{Q}$ an answer set $A(Q)$
- Each answer $W'' \in A(Q)$ restricts the feasible weights: $W'' \subset W'$
- Now, the question-answer entropy is

$$H(Y|A(Q)) = \sum_{W'' \in A(Q)} p(W'')H(Y|W'')$$

- We assume the answers $W'' \in A(Q)$ are equiprobable
- The optimal question $Q \in \mathbf{Q}$ minimizes $H(Y|A(Q))$

Stopping criterion

- Let W' be the current set of weights
- Given a set $\mathbf{W}$ of sets of weights,

$$\mathrm{IG}_{\mathsf{max}}(\mathbf{W}) = H(Y|W') - \min_{W'' \in \mathbf{W}} H(Y|W'')$$

- (relative) Stop asking questions from $\mathbf{Q}$ if

$$\mathrm{IG}_{\mathsf{max}}\Big(\bigcup_{Q \in \mathbf{Q}} \bigcup_{W' \in A(Q)} W'\Big) < \epsilon$$

- (absolute) Stop asking questions if

$$\mathrm{IG}_{\mathsf{max}}\Big(\bigcup_{w \in W'} \{w\}\Big) < \epsilon$$

Implementation

- Implementation using Monte Carlo sampling (i.e. SMAA) in R
- Sample weights with arbitrary linear constraints using HAR
  - Hit and Run presented last EURO, published in EJOR
  - Algorithm by R.L. Smith (1984)
  - Complexity proofs by L. Lovász – $O^*(n^3)$
- Often $f_X(x)$ can be sampled efficiently
  - If not, MCMC can be used, but requires convergence checking

Estimating entropy

Outcome space differs per problematic:

- Choice: m possible outcomes
- Ranking: $m!$ possible outcomes

Estimate $H(Y|A(Q))$ using sampling:

- Choice: since 10k iterations estimate rank probabilities to sufficient accuracy, the entropy is also estimated accurately
- Ranking: except in trivial problems, we are unlikely to observe all possible (or even plausible) rankings
  - However, $p(y|W') \to 0$ faster than $\log p(y|W') \to \infty$
  - Hence, it is safe to take $p(y|W') = 0$ for y not encountered
  - Is 10k samples enough in this case?

Additional restrictions

- The set of questions $\mathbf{Q}$ is finite (and reasonably small)
  - If not (e.g. pair-wise judgments), take a representative sample
- The set of answers $A(Q)$ is finite (and reasonably small)
  - If not (e.g. interval SWING), take a representative sample

Stopping criterion

- The relative criterion (based on answer entropies) is easy
  - Necessary calculations are already being done
- The absolute criterion (based on point-set entropies) is expensive
  - Requires full SMAA analysis for each point

Example: anti-thrombolytic therapy

- Two alternatives: Enoxaparin, Heparin
- Three criteria:
  - Proximal DVT (0 – 25%)
  - Distal DVT (15 – 40%)
  - Major bleeding (0 – 10%)
- Beta-distributed measurements from a clinical trial
  - Indicate a trade-off between DVT and bleeding
- We set $\mathbf{Q}$ to the corner-state alternatives
  - Equivalent to ordinal statements
- **Extremely** simple example – no time for realistic complexity

Selecting question 1

- Prox DVT ($0.25 \to 0.0$) or Dist DVT ($0.4 \to 0.15$)?
- Prox DVT ($0.25 \to 0.0$) or Bleeding ($0.1 \to 0.0$)?
- Dist DVT ($0.4 \to 0.15$) or Dist DVT ($0.1 \to 0.0$)?

Selecting question 1

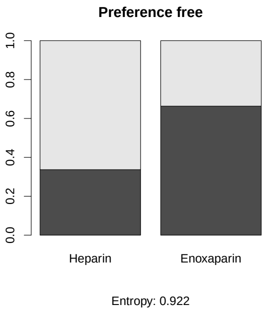

- Prox DVT (0.25 $\rightarrow$ 0.0) or Dist DVT (0.4 $\rightarrow$ 0.15)? 0.910
- Prox DVT (0.25 $\rightarrow$ 0.0) or Bleeding (0.1 $\rightarrow$ 0.0)?
- Dist DVT (0.4 $\rightarrow$ 0.15) or Dist DVT (0.1 $\rightarrow$ 0.0)?

Selecting question 1

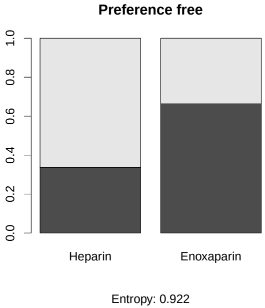

- Prox DVT ($0.25 \rightarrow 0.0$) or Dist DVT ($0.4 \rightarrow 0.15$)? 0.910
- Prox DVT ($0.25 \rightarrow 0.0$) or Bleeding ($0.1 \rightarrow 0.0$)? 0.719
- Dist DVT ($0.4 \rightarrow 0.15$) or Dist DVT ($0.1 \rightarrow 0.0$)?

Selecting question 1

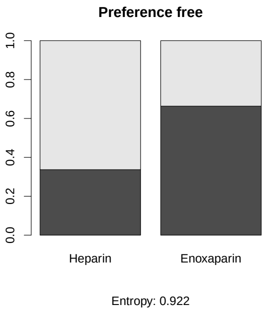

- Prox DVT ($0.25 \to 0.0$) or Dist DVT ($0.4 \to 0.15$)? 0.910
- Prox DVT ($0.25 \to 0.0$) or Bleeding ($0.1 \to 0.0$)? 0.719
- Dist DVT ($0.4 \to 0.15$) or Dist DVT ($0.1 \to 0.0$)? 0.816

Selecting question 1

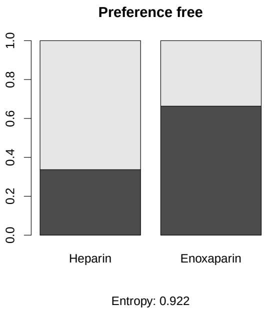

- Prox DVT ($0.25 \rightarrow 0.0$) or Dist DVT ($0.4 \rightarrow 0.15$)? 0.910
- Prox DVT ($0.25 \rightarrow 0.0$) or Bleeding ($0.1 \rightarrow 0.0$)? 0.719
- Dist DVT ($0.4 \rightarrow 0.15$) or Dist DVT ($0.1 \rightarrow 0.0$)? 0.816

Selecting question 2

- Prox DVT (0.25 $\to$ 0.0) or Dist DVT (0.4 $\to$ 0.15)?
- Prox DVT (0.25 $\to$ 0.0) or Bleeding (0.1 $\to$ 0.0)?
- Dist DVT (0.4 $\to$ 0.15) or Dist DVT (0.1 $\to$ 0.0)?

Selecting question 3

- Prox DVT ($0.25 \to 0.0$) or Dist DVT ($0.4 \to 0.15$)?
- Prox DVT ($0.25 \to 0.0$) or Bleeding ($0.1 \to 0.0$)?
- Dist DVT ($0.4 \to 0.15$) or Dist DVT ($0.1 \to 0.0$)?

Result

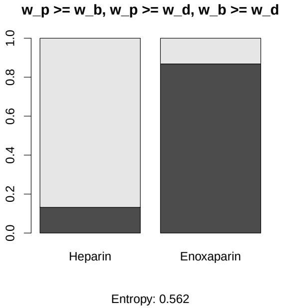

- Prox DVT $(0.25 \to 0.0)$ or Dist DVT $(0.4 \to 0.15)$?
- Prox DVT $(0.25 \to 0.0)$ or Bleeding $(0.1 \to 0.0)$?
- Dist DVT $(0.4 \to 0.15)$ or Dist DVT $(0.1 \to 0.0)$?

Entropy evolution

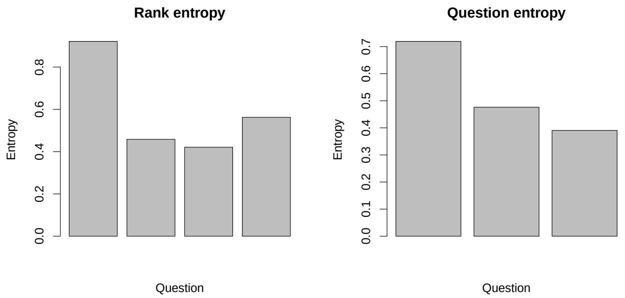

Entropy evolution

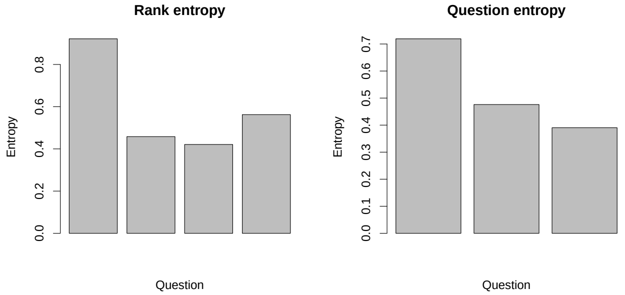

- Actual entropy increased after Q3
- Q3 conditional entropy very close to entropy after Q2 (~ 0.4)
- However, answer entropies were 0.562 and 0.211

Global stopping criterion

- Range evolved from $(0.05, 1.0)$ to $(0.09, 0.95)$
  - Probably due to not ruling out corners of pref. space
- Hence, *other* preference information desirable
- We elicited imprecise trade-off ratios from DM
  - Entropy range reduced to $(0.437, 0.529)$
  - Is this ($\epsilon \approx 0.09$) sufficiently narrow to stop?

Discussion

- The $W'' \in A(Q)$ are assumed equiprobable
  - But, we already have f_W, so why not $p(W'') = \int_{w \in W_n} f_W(w \in W'' | w \in W')\mathrm{d}w$?
- How do we estimate the absolute stopping criterion?
  - How many points should be sampled to get a robust minimum?
  - Should we use a quantile (e.g. 0.05) instead?
  - Alternatives' central weights instead of a random sample?
- How do we set ϵ?