Decision Methods and Models

Master's Degree in Computer Science

Roberto Cordone
DI - Università degli Studi di Milano

Schedule:	Thursday 16.30 - 18.30 in Aula Magna (CS department)
	Friday 12.30 - 14.30 in classroom 301
Office hours:	on appointment
E-mail:	roberto.cordone@unimi.it
Web page:	https://homes.di.unimi.it/cordone/courses/2024-mmd/2024-mmd.html
Ariel site:	https://myariel.unimi.it/course/view.php?id=4467

Weak orders with inconsistent weights

We assume

- a preference relation Π that is a weak order with an additive utility
- a certain environment: $|\Omega| = 1 \Rightarrow f(x, \bar{\omega})$ reduces to $f(x)$
- a single decision-maker: $|D| = 1 \Rightarrow \Pi_d$ reduces to Π

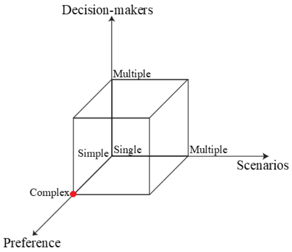

However

- the pairwise comparison matrix $\tilde{\Lambda}$ is incorrect
- the normalised utilities $\tilde{u}(f_l)$ are incorrect

$$\textit{Therefore, } u(f) = \sum_{l \in P} w_l \tilde{u}_l(f_l) \textit{ makes no sense}$$

Criticisms to the Multi-Attribute Utility Theory

The process to build $u(f)$ is tedious, complex and error-prone

Saaty (1980) remarks that in practice

1. the normalised utilities $\tilde{u}_l(x)$ are only approximated

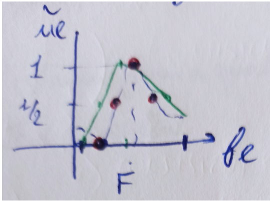

2. the pairwise comparison coefficients $\tilde{\lambda}_{lm}$ are only approximated: $\tilde{\Lambda}$ is typically positive and reciprocal, but inconsistent matrices give nonunivocal weight vectors

$$\begin{bmatrix} 1 & 2 & 4 \\ 1/2 & 1 & 3 \\ 1/4 & 1/3 & 1 \end{bmatrix}$$

3. the approximations of w_l and $\tilde{u}_l(x)$ combine in cascade

$$\tilde{u}(x) = \sum_{l \in P} (w_l \pm \delta w_l)(\tilde{u}_l(x) \pm \delta \tilde{u}_l(x))$$

Correcting the pairwise comparison matrix

A possible approach is force consistency modifying $\tilde{\Lambda}$ as little as possible

This requires a definition of distance between matrices

Given a metric $d : \mathbb{R}^{p^2} \times \mathbb{R}^{p^2} \rightarrow \mathbb{R}$, we can solve the problem

$$\min d\left(W, \tilde{\Lambda}\right)$$

$$W_{lm} = \frac{w_l}{w_m}$$

$$w_l > 0 \qquad l \in P$$

$$\sum_{l \in P} w_l = 1$$

to determine an unknown matrix W that is

- positive, reciprocal and consistent
- as close as possible to the known matrix $\tilde{\Lambda}$

Metrics on matrices

But how is the distance between two matrices defined?

Infinite different definitions are possible (e.g., the Lebesgue distances)

- Manhattan distance L_1:

$$d_1\left(W,\tilde{\Lambda}\right) = \sum_{l,m\in P}\left|\frac{w_l}{w_m} - \tilde{\lambda}_{lm}\right|$$

- Euclidean distance L_2:

$$d_2\left(W,\tilde{\Lambda}\right) = \sqrt{\sum_{l,m\in P}\left(\frac{w_l}{w_m} - \tilde{\lambda}_{lm}\right)^2}$$

- L_∞ distance:

$$d_\infty\left(W,\tilde{\Lambda}\right) = \max_{l,m\in P}\left|\frac{w_l}{w_m} - \tilde{\lambda}_{lm}\right|$$

The optimal matrix W° depends on the distance chosen, that is arbitrary

Eigenvalues and consistency

Another way to achieve consistency exploits the following properties

Given a square matrix $\tilde{\Lambda}$ of order $p \times p$

- eigenvalues are the p solutions of equation $|\lambda I - \tilde{\Lambda}|$
- eigenvectors associated to λ are the ∞ nonzero solutions of $\lambda x = \tilde{\Lambda} x$

If $\tilde{\Lambda}$ is positive, reciprocal and consistent

1. the dominant (maximum absolute value) eigenvalue is $\lambda_{max} = p$
2. the other $p - 1$ eigenvalues are equal to zero ($\tilde{\Lambda}$ has rank 1)
3. the eigenvectors associated to λ_{max} are proportional to the weight vector, and w is the normalisation of any dominant eigenvector x_{max}

The eigenvalue method

Consequently, the eigenvalue method proposes to

- compute the eigenvalues and identify the dominant one, $\lambda_{\max}$
- compute the associated dominant eigenvector $x_{\max}$
- normalise it to obtain the weight vector $w = \dfrac{x_{\max}}{\|x_{\max}\|}$
- build the corrected matrix as $W = \left\{ \dfrac{w_l}{w_m} \right\}$

Problem: the resulting matrix W can be far from $\tilde{\Lambda}$

Other consistent matrices can be closer to $\tilde{\Lambda}$, even elementwise!

A counter-proposal

Can a forcedly consistent matrix be safely used? The debate is sharp

An alternative approach is to

- consciously accept imprecise values for w_l and $\tilde{u}_l(x)$
- compute them based on the stronger aspects of human psychology
- aim at a qualitative ranking, instead of a quantitative one

For example, the Analytic Hierarchy Process (*AHP*) by Saaty (1980)

- replaces absolute measures with relative ones
- replaces quantitative ratios with qualitative scales
- builds a hierarchy of indicators in order to
  compare only conceptually similar quantities

Relative measures

Psychology suggests that humans are bad at absolute judgments, but they are better at relative ones

Idea: replace absolute estimates of utility with pairwise comparisons

Instead of building $|X|$ absolute utility estimates $\tilde{u}_l(x)$ for each $x \in X$

- build $|X|^2$ utility ratio estimates

$$\lambda^l_{xy} \approx \frac{\tilde{u}_l(x)}{\tilde{u}_l(y)} \quad \text{for each } x, y \in X$$

- make them consistent with the methods discussed above (if necessary)
- derive pseudoutilities $\tilde{u}_{lx}$ from the consistent matrix

Qualitative scales

Psychology: humans are unable to distinguish more than 5 levels

Idea: replace quantitative ratios with qualitative ones

Instead of guessing $\frac{\tilde{u}_l(x)}{\tilde{u}_l(y)}$ for each $x, y \in X$

choose a representative value in Saaty's scale

- 1 for equally good
- 3 for moderately better and $1/3$ for moderately worse
- 5 for strongly better and $1/5$ for strongly worse
- 7 for very strongly better and $1/7$ for very strongly worse
- 9 for absolutely better and $1/9$ for absolutely worse

Do the same for the weight ratios $\tilde{\lambda}_{lm} = \frac{w_l}{w_m}$ for each $l, m \in P$

- choose representative qualitative values λ_{lm}
- impose consistency
- derive pseudoweights w_l from the consistent matrix

Indicator hierarchy

Psychology: humans are bad at comparing dishomogeneous things

Idea: build an indicator tree and compare only sibling indicators

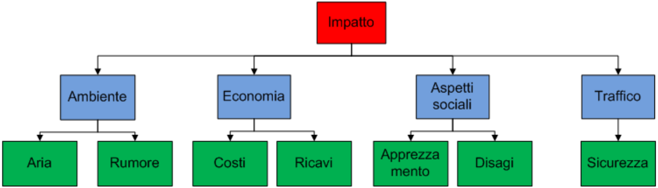

Instead of comparing l and m for each $l, m \in P$

- estimate weight ratio λ_{lm} only when l and m have the same father
- perform comparisons at all levels of the indicator tree, that is not only between indicators, but also between indicator groups

Replace a single big pairwise comparison matrix with many small ones

Overall computation

Given the pseudoutilities $\tilde{u}_{lx}$ and the pseudoweights w_l

- combine pseudoweights and pseudoutilities with convex combinations
- combine the pseudoweights with products from root to leaf
  this corresponds to normalising each set of sibling nodes

$$u(x) = \sum_{l \in P} \prod_{\ell \in \gamma_l} w_\ell \tilde{u}_{lx}$$

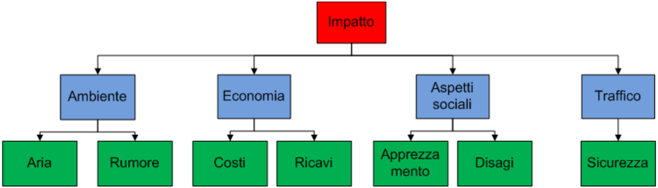

Be aware that the resulting numbers are approximations

- handle with care
- perform sensitivity analyses

Example

See example 56 on the lecture notes to discuss

- pairwise comparisons to derive pseudoutilities
- pairwise comparisons to derive pseudoweights
- the convex combination of pseudoutilities and pseudoweights

We assume consistency for the sake of simplicity (and for the small size)

Notice that

- pseudoutilities always sum to 1,
  normalised utilities generally do not
- pseudoutilities cannot be equal to 0 or 1,
  normalised utilities can
- normalised w_l and u_{lx} imply a normalised u_x

Example

See example 57 on the lecture notes

- here we use absolute utilities provided by reliable experts

  It is a hybrid approach
- we can reduce the tree to a single level multiplying the weights root to leaf
  or make backward induction on each alternative

Rank reversal

A practical decision process goes through loops of alternative identification

- new alternatives can be generated in different phases

This obviously modifies the feasible region X: new best solutions could appear

In the *AHP* adding or removing alternatives can change the ranking of the other alternatives

This is counterintuitive and does not occur in the classical *MAUT*

The reason is that

- the absolute utilty $\tilde{u}_l(x)$ depends only on x
- the pseudoutilty u_{xl} depends on all $x \in X$

Example

See in example 57 on the lecture notes the introduction of a_4, that is

- intermediate for price
- indifferent to a_2 and a_3 for size
- good for zone
- indifferent to a_3 for purchase conditions

Through the normalisation, a_4 subtracts pseudoutility

- strongly from a_3 for zone and conditions, where a_3 is good
- weakly from a_1 for price and size, where a_1 is bad

Therefore $a_3 \prec a_2 \prec a_1$ reverses into $a_1 \prec a_2 \prec a_3$

Rank reversal

A result depending on X is undesirable, because

- X is not always given a priori
- modifying X allows to manipulate the result

Unfortunately, all decision processes based on pairwise comparisons between alternatives suffer from rank reversal

- sport tournaments
- elections
- ...

How to avoid rank reversal?

Avoiding rank reversal

Rank reversal can be avoided using absolute scales and a priori estimates:

- fix a (finite!) set of absolute levels for each indicator
- make pairwise comparisons on levels, instead of alternatives
- evaluate each alternative assigning it to a level for each indicator

In this way, the pseudoutilities refer to absolute levels, fixed once for all, and not depending on the alternatives

Example

See example 58 on the lecture notes

Advantages and disadvantages

The use of absolute classes also allows

- open decision processes, in which the alternatives arrive gradually
- very long decision processes, in which alternatives arriving at far away times cannot be compared significantly (e.g., hiring and recruiting)
- to end the process as soon as an alternative reaches a satisfactory threshold (this limits the cost of the decision process)

However, the trick also introduces further approximations

- rather different values can be flattened putting them into the same class
- very similar values can be strongly differentiated putting them into separate classes