# Preference Handling

Alexis Tsoukiàs[1] and Paolo Viappiani[2]

1) LAMSADE - CNRS, Université Paris-Dauphine
www.lamsade.dauphine.fr/∼tsoukias
2) LIP6 - CNRS, Université Pierre et Marie Curie
www-desir.lip6.fr/∼viappianip/

ACM RecSys Conference, Hong Kong

October 13, 2013

# Outline

# Aims of the Tutorial

- Introducing the general field that deals with techniques for representing, learning and reasoning with preferences
- Community that studies preference from a formal algorithmic point of view: *Algorithmic Decision Theory*
- The tutorial will introduce both the formalisms more widely used to model and represent preferences as well as the procedures aimed at learning preferences and at producing a recommendation for an end-user (decision-maker)

# Preference Handling Systems are Everywhere

- Not only recommender systems
- Computational advertisement
- Intelligent user interfaces
- Cognitive assistants
- Personalized medecine
- Personal Robots

*What the theory has to say about preferences?*

# Outline

# What are Preferences?

- Preferences are “rational” desires.
- Preferences are at the basis of any decision aiding activity.
- There are no decisions without preferences.
- Values, Likelihoods, Opinions, ...
- ... but also Objectives, Desires, Utilities, Beliefs,...

# Preference Statements

- I like red shoes.
- I do not like brown sugar.
- I prefer Obama to McCain.
- I do not want tea with milk.
- Cost is more important than safety.
- I prefer flying to Athens than having a suite at Istanbul.

# Preference Statements

- I like red shoes.
- I do not like brown sugar.
- I prefer Obama to McCain.
- I do not want tea with milk.
- Cost is more important than safety.
- I prefer flying to Athens than having a suite at Istanbul.

**Three issues:**

- Relative vs. Absolute assessments
- Single vs Multi-attribute assessments
- Positive vs Negative assessments

# Binary relations

## Preference Relation

$\succeq \; \subseteq \; A \times A$ or $\succeq \; \subseteq \; A \times P \cup P \times A$

$x \succeq y$ stands for $x$ is at least as good as $y$

$A$ is the set on which we express our preferences
$P$ is the set (if necessary) of norms, standards, references to which we may compare elements of $A$
At this stage we do not care if $A$ or $P$ are described under multiple attributes.

# Symmetry and Asymmetry

$\succeq$ can be decomposed into an asymmetric and a symmetric part

## Asymmetric part

- Strict preference $\succ$: $x \succeq y \ \wedge \ \neg(y \succeq x)$

## Symmetric part

- Indifference $\sim_1$: $x \succeq y \ \wedge \ y \succeq x$
- Incomparability $\sim_2$: $\neg(x \succeq y) \ \wedge \ \neg(y \succeq x)$

# Binary relations properties

Binary relations have specific properties such as:

- Irreflexive: $\forall x \ \neg(x \succ x)$;
- Asymmetric: $\forall x, y \ x \succ y \rightarrow \neg(y \succ x)$;
- Transitive: $\forall x, y, z \ x \succ y \wedge y \succ z \rightarrow x \succ z$;
- Ferrers; $\forall x, y, z, w \ x \succ y \wedge z \succ w \rightarrow x \succ w \vee z \succ y$;

# $\sim_1, \sim_2, \succ$ Preference Structures

**If $\succeq$ is a weak order then:**

$\succ$ is a strict partial order, $\sim_1$ is an equivalence relation and $\sim_2$ is empty.

# $\sim_1, \sim_2, \succ$ Preference Structures

**If $\succeq$ is a weak order then:**

$\succ$ is a strict partial order, $\sim_1$ is an equivalence relation and $\sim_2$ is empty.

**If $\succeq$ is an interval order then:**

$\succ$ is a partial order of dimension two, $\sim_1$ is not transitive and $\sim_2$ is empty.

# What characterises such structures?

## Characteristic Properties

Weak Orders are complete and transitive relations.
Interval Orders are complete and Ferrers relations.

# What characterises such structures?

## Characteristic Properties

Weak Orders are complete and transitive relations.
Interval Orders are complete and Ferrers relations.

## Why this is important?

Because it allows to establish *representation theorems* which tell us under which conditions these preference structures have a numerical representation.

# Numerical Representations

## Weak Orders

If $\succeq$ is a w.o. $\Leftrightarrow \exists f : A \mapsto \mathbb{R} : x \succeq y \leftrightarrow f(x) \geq f(y)$

# Numerical Representations

## Weak Orders

If $\succeq$ is a w.o. $\Leftrightarrow$ $\exists f : A \mapsto \mathbb{R} : x \succeq y \leftrightarrow f(x) \geq f(y)$

## Interval Orders

If $\succeq$ is an i.o.
$\Leftrightarrow$ $\exists f, g : A \mapsto \mathbb{R} : f(x) > g(x); x \succeq y \leftrightarrow f(x) \geq g(y)$

# What if $A$ is multi-attribute described?

$$x = \langle x_1 \cdots x_n \rangle \quad y = \langle y_1 \cdots y_n \rangle$$

# What if $A$ is multi-attribute described?

$$x = \langle x_1 \cdots x_n \rangle \quad y = \langle y_1 \cdots y_n \rangle$$

$$x \succeq y \quad \Leftrightarrow \quad \Phi([u_1(x_1) \cdots u_n(_n)], [u_1(y_1) \cdots u_n(y_n)] \geq 0$$

# Example

Suppose you have 4 projects $x, y, z, w$ of urban rehabilitation and an assessment dimension named “land use”. You have:

- $d_l(x) = 100$sqm;
- $d_l(y) = 50$sqm;
- $d_l(z) = 1000$sqm;
- $d_l(w) = 500$sqm;

Preferences expressed could be for instance (suppose the decision maker dislikes land use):

$d_l(y) \succ d_l(x) \succ d_l(w) \succ d_l(z)$

A possible numerical representation could thus be:

$h_l(y) = 4$, $h_l(x) = 3$, $h_l(w) = 2$ $h_l(z) = 1$, but also:

$h_l(y) = 50$, $h_l(x) = 100$, $h_l(w) = 500$, $h_l(z) = 1000$

# Is this sufficient?

**NOT always!**

# Is this sufficient?

## NOT always!

We may need something more rich. We may need to know, when we compare $x$ to $y$ (and we prefer $x$) if this preference is “stronger” to the one expressed when comparing (on the same dimension) $z$ fo $w$.

**We need to compare differences of preferences**

# A Value function

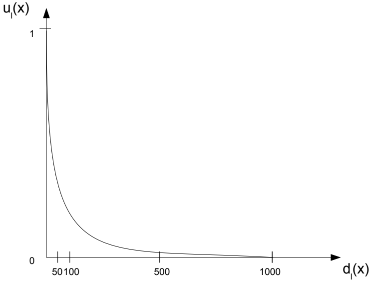


For instance, if the above function represents the value of "land use" it is clear that the difference between 50sqm and 100sqm is far more important from the one between 500sqm and 1000sqm.

# Is all that sufficient?

**NOT always!**

If $A$ is described on multiple attributes

# Is all that sufficient?

**NOT always!**

If *A* is described on multiple attributes

1. The problem is that we need to be able to compare the differences of preferences on one dimension to the differences of preferences on another one (let's say differences of preferences on land use with differences of preferences on esthetics.
2. At the same time we need to take into account the intuitive idea that for a given decision maker certain dimensions are more "important" than other ones.

# Principal Tests

1. The different dimensions are separable.
2. Preferences on each dimension are independent.
3. Preferences on each dimension are measurable in terms of differences.

# Outline

1. What it is?
2. Using Preferences
   - Problem Statements
   - Preference Aggregation
   - The Logic of Preferences
   - Rules
   - Recommendations
3. Learning Preferences
4. A little bit further

# What is a decision problem?

Consider a set $A$ established as any among the following:

- an enumeration of objects;
- a set of combinations of binary variables (possibly the whole space of combinations);
- a set of profiles within a multi-attribute space (possibly the whole space);
- a vector space in $\mathbb{R}^n$.

# What is a decision problem?

Consider a set $A$ established as any among the following:

- an enumeration of objects;
- a set of combinations of binary variables (possibly the whole space of combinations);
- a set of profiles within a multi-attribute space (possibly the whole space);
- a vector space in $\mathbb{R}^n$.

**Technically:**

A Decision Problem is a partitioning of $A$ under some desired properties.

# Partitioning? For what?

Practically we partition *A* in *n* classes. These can be:

| | Pre-defined wrt some external norm | Defined only through relative comparisons |
|---|---|---|
| Ordered | | |
| Not Ordered | | |

# Partitioning? For what?

Practically we partition *A* in *n* classes. These can be:

| | Pre-defined wrt some external norm | Defined only through relative comparisons |
|---|---|---|
| Ordered | Rating | Ranking |
| Not Ordered | Assigning | Clustering |

# Partitioning? For what?

Practically we partition *A* in *n* classes. These can be:

| | Pre-defined wrt some external norm | Defined only through relative comparisons |
|---|---|---|
| Ordered | Rating | Ranking |
| Not Ordered | Assigning | Clustering |

Two special cases:
- there are only two classes (thus complementary);
- the size (cardinality) of the classes is also predefined.

# The Problem

Suppose we have $n$ ordering relations $\succeq_1 \cdots \succeq_n$ on the set $A$.
We are looking for an overall ordering relation $\succeq$ on $A$ "representing" the different orders.

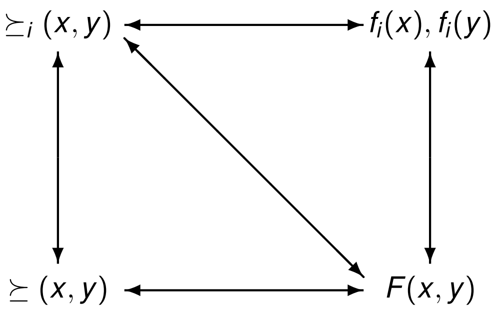

# How do we do that?

**The Borda path**

Give a value to the rank of each element of $x \in A$ and then sum the ranks. Then compare the sum of the ranks.

# How do we do that?

## The Borda path

Give a value to the rank of each element of $x \in A$ and then sum the ranks. Then compare the sum of the ranks.

## The Condorcet path

Compare each $x \in A$ to all other elements $y \in A$, then use majority to establish who is better among $x$ and $y$.

# Generalising Borda

$$x \succeq y \;\Leftrightarrow\; \mathcal{F}(u_j(x), r_j(y)) \geq 0$$

# Generalising Borda

$$x \succeq y \Leftrightarrow \mathcal{F}(u_j(x), r_j(y)) \geq 0$$

Special case:

$$x \succeq y \Leftrightarrow \sum_j u_j(x), \geq \sum_j u_j(y)$$

What do we need to know?

# Generalising Borda

$$x \succeq y \;\Leftrightarrow\; \mathcal{F}(u_j(x), r_j(y)) \geq 0$$

Special case:

$$x \succeq y \;\Leftrightarrow\; \sum_j u_j(x), \geq \sum_j u_j(y)$$

What do we need to know?

the primitives: $\succeq_j \;\subseteq\; A \times A$
Differences of preferences:
- $(xy)_1 \succcurlyeq (zw)_1$
- $(xy)_1 \succcurlyeq (zw)_2$

# Generalising Condorcet

$$x \succeq y \;\Leftrightarrow\; H_{xy} \geq H_{yx}$$

What do we need to know?

# Generalising Condorcet

$$x \succeq y \ \Leftrightarrow \ H_{xy} \geq H_{yx}$$

What do we need to know?

the primitives: $\succeq_j \ \subseteq \ A \times A$
An ordering relation on $2^{\succeq_j}$

# Advantages and drawbacks

## The Borda path

You have values. If correctly adopted, you have real values (meaningful measures of differences of preferences. Easy to bias. Easy to use it making big mistakes ... Too much axioms to satisfy.

# Advantages and drawbacks

## The Borda path

You have values. If correctly adopted, you have real values (meaningful measures of differences of preferences. Easy to bias. Easy to use it making big mistakes ... Too much axioms to satisfy.

## The Condorcet path

You only need the ordinal pairwise comparisons. No need to measure quantitatively the differences of preferences. Less axioms to satisfy. The drawback is that you need further steps to become operational

# Representation Problem

Often impossible to state explicitly the preference relation ($\succeq$), especially when A has a combinatorial structure:

*"I prefer flying than driving to Paris"* could then be:

$$p \wedge f \wedge \neg d \succ p \wedge \neg f \wedge d$$

# Representation Problem

Often impossible to state explicitly the preference relation ($\succeq$), especially when A has a combinatorial structure:

*"I prefer flying than driving to Paris"* could then be:

$$p \wedge f \wedge \neg d \succ p \wedge \neg f \wedge d$$

**What is the advantage?**

You can use SAT and more generally constraint satisfaction algorithms in order to "solve" the preference aggregation problem including complex conditional preference statements

# Different tools

## Logical Languages

- Weighted Logics
- Conditional Logics
- ...

# Different tools

## Logical Languages

- Weighted Logics
- Conditional Logics
- ...

## Graphical Languages

- Conditional Preference networks (CP nets)
- Conditional Preference networks with trade-offs
- Generalised Additive Independence networks

# Association and dominance

## Association rules

$$\langle x_1 = \kappa, x_2 = \mu \cdots x_n = \lambda \rangle \rightarrow x \in K_l$$

$$\langle x_1 \geq \kappa, x_2 \geq \mu \cdots x_n \geq \lambda \rangle \rightarrow x \in K_l$$

# Association and dominance

## Association rules

$$\langle x_1 = \kappa, x_2 = \mu \cdots x_n = \lambda \rangle \ \rightarrow \ x \in K_l$$

$$\langle x_1 \geq \kappa, x_2 \geq \mu \cdots x_n \geq \lambda \rangle \ \rightarrow \ x \in K_l$$

## Dominance rules

$$\langle x_1 \succeq^{\kappa} y_1, x_2 \succeq^{\mu} y_2 \cdots x_n \succeq^{\lambda} y_n \rangle \ \rightarrow \ x \succeq^{k} y$$

# Advantages and drawbacks

- Rules fit well for rating and assignment problems, less for other ones.
- Asymmetric models are easier to work with.
- Any logical or rule based model has an equivalent conjoint measurement model leading either to the Borda or the Condorcet path.

# General

Preference relations do not automatically lead to operational recommendations

*The Borda path is more likely to do it, but is expensive and manipulable*

*The Condorcet path is more likely to need further elaborations, but is less expensive and more robust*

# Lessons learned

1. There is no universal preference aggregation procedure and there will never exist one.
2. The choice of an aggregation procedure is part of the modelling activity and need to be justified.
3. Knowing the axiomatics behind aggregation allows to make reasoned choices.

# Claim

## Choosing is framed!

The properties of the set $A$, the type of problem statement, the holding or not of preferential independence, the explicit use of differences of preferences, the explicit use of negative preference statements, are the necessary and sufficient features for choosing, designing, justifying and axiomatically characterising any decision problem and the associated resolution methods and algorithms.

# Outline

# Why do we want to learn?

- Explaining
- Understanding
- Justifying
- Prescribing, Recommending

# What are we trying to learn?

- Preferences ($x$ is better than $y$);
- Second Order Preferences (cost is more important than mass);
- Models (observed decision behaviours);
- Parameters (of a fixed decision model)

# Preferences

- Learn if $x \succeq y$.
- Test if $\succeq$ satisfies any nice properties.
- Test of $\succeq$ has a numerical representation and construct it.
- What is the minimal number of questions, given a set $A$?

# Second Order Preferences

**Example**

If $x \succeq_1 y$ and $y \succeq_2 x$ and $x \succeq y$, then $\succeq_1 \rhd \succeq_2$ (this applies for values, opinions and likelihoods).

# Second Order Preferences

## Example

If $x \succeq_1 y$ and $y \succeq_2 x$ and $x \succeq y$, then $\succeq_1 \rhd \succeq_2$ (this applies for values, opinions and likelihoods).

## Weights ...

Weights are everywhere, but they do not exist ... they are powers of coalitions, trade-offs, etc..

# Examples

## Standard weighting parameters (weights attached to criteria)

- **Weighted sum** : $u(\mathbf{x};\omega) = \sum_{i=1}^{n} \omega_i x_i$
- **Weighted Tchebycheff** : $u(\mathbf{x};\omega) = \max_{i \in [\![1;n]\!]} \{\omega_i \frac{x_i^* - x_i}{x_i^* - x_{*i}}\}$

## Rank-dependent weighting parameters (weights attached to ranks)

- **OWA** : $u(\mathbf{x};w) = \sum_{i=1}^{n} w_i x_{(i)}$
- **Choquet** : $u(\mathbf{x};v) = \sum_{i=1}^{n} \left[x_{(i)} - x_{(i-1)}\right] v(X_{(i)})$

# How do we learn preferences?

1. Directly (asking...)
2. Indirectly
    - Examples, Cases
    - Global Preference Statements
    - Observing user's clicking behavior

# Classic Approaches for Utility Elicitation

- Assessment of multi attribute utility functions
  - Typically long list of questions
  - Focus on high risk decision
  - Goal: learn utility parameters (weights) up to a small error
- Which queries?
  - Local: focus on attributes in isolation
  - Global: compare complete outcomes
- Standard Gamble Queries
  - Choose between option $x_0$ for sure or a gamble $< x^{\top}, l, x_{\perp}, 1-l >$ (best option $x^{\top}$ with probability $l$, worst option $x_{\perp}$ with probability $1-l$)

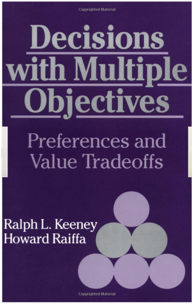

# Standard Elicitation: Additive Models

- Consider an attribute (for example, color)
  - Ask for the best value (say, red)
  - Ask for worst value (gray)
  - Ask *local standard gamble* for each remaining color to assess it local utility value (*value function*)
- Refine intervals on local utility values
  - Bound queries
- Scaling factors
  - Define reference outcome
  - Ask global queries in order to assess the difference in utility occurring when, starting from the reference outcome, “moving” a particular attribute to the best / worst

# Automated Elicitation vs Classic Elicitation

## Problems with the classic view

- Standard gamble queries (and similar queries) are difficult to respond
- Large number of parameters to assess
- Unreasonable precision required
- Cognitive or computational cost may outweigh benefit

## Automated Elicitation and Recommendation

Important points:

- Cognitively plausible forms of interaction
- Incremental elicitation until a decision is possible
- We can often make optimal decisions without full utility information
- Generalization across users

# Adaptive Utility Elicitation

**Utility-based Interactive Recommender System:**

- *Bel*: the system's "belief" about the user's utility function $u$
- $Opt(Bel)$: optimal decision given incomplete beliefs about $u$

## Algorithm: Adaptive Utility Elicitation

1. **Repeat** until *Bel* meets some termination condition
    1. Ask user some query
    2. Observe user response $r$
    3. Update *Bel* given $r$
2. Recommend $Opt(Bel)$

## Types of Beliefs

- *Probabilistic Uncertainty:* distribution of parameters, updated using Bayes
- *Strict Uncertainty:* feasible region (if linear constraints: convex polytope)

# Minimax Regret

## Intuition

Adversarial game; the recommender selects the item reducing the “regret” wrt the “best” item when the uniknown parameters are chosen by the adversary

- Robust criterion for decision making under uncertainty [Savage; Kouvelis]
- Effective for decision and elicitation under utility uncertainty [Boutilier et al., 2006]

## Advantages

- Easy to update our knowledge about the user: whenever a new preferences is added, we just restrict more the feasible region (polytope)
- No “prior” assumption required
- MMR computation suggests queries to ask to the user

## Limitations

- No account for noisy responses
- Formulation of the optimization depends on the assumption about the utility

# Minimax Regret

**Assumption:** a set of feasible utility functions $W$ is given

**The *pairwise max regret***

$$PMR(\mathbf{x}, \mathbf{y}; W) = \max_{w \in W} \; u(\mathbf{y}; w) - u(\mathbf{x}; w)$$

**The *max regret***

$$MR(\mathbf{x}; W) = \max_{\mathbf{y} \in \mathbf{X}} PMR(\mathbf{x}, \mathbf{y}; W)$$

**The *minimax regret***

$MMR(W)$ of $W$ and the *minimax optimal item* $\mathbf{x}^*_W$:

$$\begin{aligned} MMR(W) &= \min_{\mathbf{x} \in \mathbf{X}} MR(\mathbf{x}, W) \\ \mathbf{x}^*_W &= \arg\min_{\mathbf{x} \in \mathbf{X}} MR(\mathbf{x}, W) \end{aligned}$$

# Example

| *item* | *feature1* | *feature2* |
|---|---|---|
| a | 10 | 14 |
| b | 8 | 12 |
| c | 7 | 16 |
| d | 14 | 9 |
| e | 15 | 6 |
| f | 16 | 0 |

Linear utility model with normalized utility weights ($w_1 + w_2 = 1$);
$u(x; w) = (1 - w_2)x_1 + w_2 x_2 = (x_2 - x_1)w_2 + x_1$

Notice: it is a 1 dimensional problem

Initially, we only know that $w_2 \in [0, 1]$

$PMR(a, f; w_2) = \max_{w_2} u(f; w_2) - u(a; w_2)$
$= \max_{w_2} 6(1 - w_2) - 14w_2 = \max_{w_2} 6 - 20w_2$
$= 6$ (for $w_2 = 0$)

$PMR(a, b; w_2) = \max_{w_2} u(b; w_2) - u(a; w_2) < 0$
(a dominates b; there can't be regret in choosing a instead of b!)

$PMR(a, c; w_2) = \max_{w_2} -3(1 - w_2) - 2w_2 = 2$
(for $w_2 = 1$)

....

# Example (continued)

| item | feature1 | feature2 |
|---|---|---|
| a | 10 | 14 |
| b | 8 | 12 |
| c | 7 | 16 |
| d | 14 | 9 |
| e | 15 | 6 |
| f | 16 | 0 |

Linear utility model with normalized utility weights ($w_1 + w_2 = 1$);
$u(x; w) = (1 - w_2)x_1 + w_2 x_2 = (x_2 - x_1)w_2 + x_1$

Notice: it is a 1 dimensional problem

Computation of the pairwise regret table.

| $PMR(\cdot, \cdot)$ | *a* | *b* | *c* | *d* | *e* | *f* | *MR* |
|---|---|---|---|---|---|---|---|
| a | 0 | -2 | 2 | 4 | 5 | **6** | **6** |
| b | 2 | 0 | 4 | 6 | 7 | **8** | 8 |
| c | 3 | 1 | 0 | 7 | 8 | **9** | 9 |
| d | 5 | 3 | **7** | 0 | 1 | 2 | 7 |
| e | 8 | 6 | **10** | 3 | 0 | 1 | 10 |
| f | 14 | 12 | **16** | 9 | 6 | 0 | 16 |

The MMR-optimal solution is *a*, adversarial choice is *f*, and minimax regret value is 6.

In reality no need to compute the full table (tree search methods) [Braziunas, PhD Thesis, 2011]

Now, we want to ask a new query to improve the decision. A very successful strategy (thought generally not optimal!) is the *current solution strategy*: ask user to compare *a* and *f*

# A graphical illustration for linear utility model (1/3)

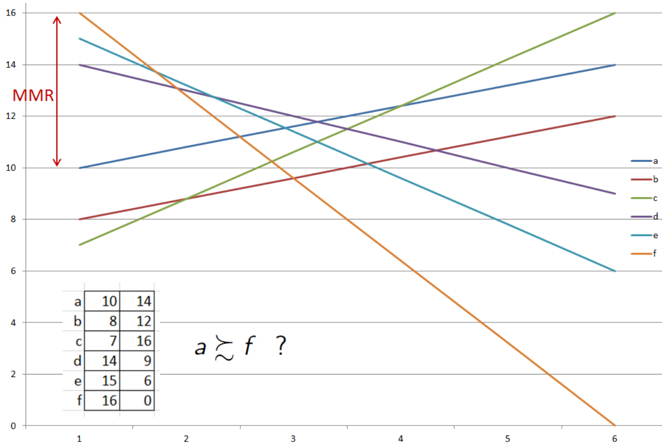


| | | |
|---|---|---|
| a | 10 | 14 |
| b | 8 | 12 |
| c | 7 | 16 |
| d | 14 | 9 |
| e | 15 | 6 |
| f | 16 | 0 |

$a \succsim f$ ?

# A graphical illustration for linear utility model (2/3)

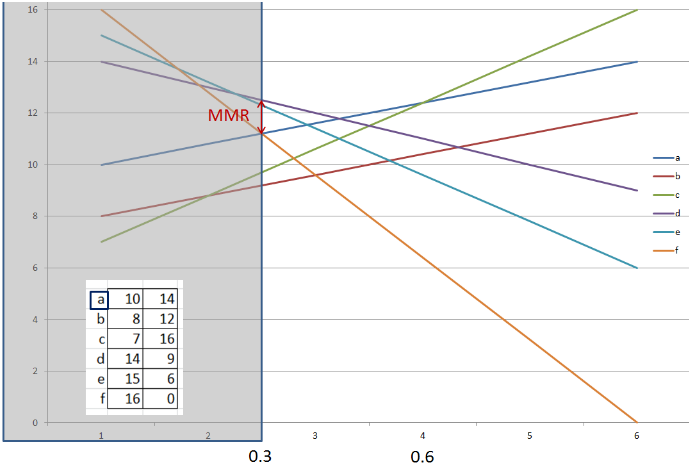


| | | |
|---|---|---|
| a | 10 | 14 |
| b | 8 | 12 |
| c | 7 | 16 |
| d | 14 | 9 |
| e | 15 | 6 |
| f | 16 | 0 |

# A graphical illustration for linear utility model (3/3)

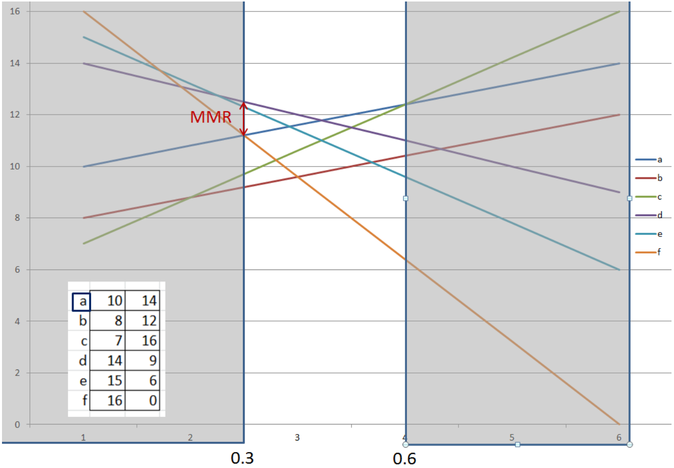


| | | |
|---|---|---|
| a | 10 | 14 |
| b | 8 | 12 |
| c | 7 | 16 |
| d | 14 | 9 |
| e | 15 | 6 |
| f | 16 | 0 |

# Bayesian Framework for Recommendation and Elicitation

- Let's assume utility $u(x; w)$ parametric in $w$ for a given structure, for example $u(x; w) = w \cdot x$
- $P(w)$ probability distribution over utility function
- Expected utility of a given item $x$
  $EU(x) = \int u(x) P(w) dw$
- Current expected utility of best recommendation $x^*$
  $EU^* = \max_{x \in A} EU(x); x^* = \arg\max_{x \in A} EU(x)$
- When a new preference is known (for instance, user prefers apples over orange), the distribution is updated according to Bayes (Monte Carlo methods, Expectation Propagation)

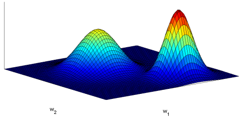


*(possible prior distribution)*

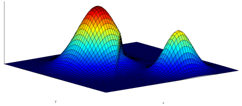

*(distribution updated after user feedback)*

# Outline

1. What it is?
2. Using Preferences
3. Learning Preferences
4. A little bit further
   - Algorithmics
   - Argumentation
   - Research
   - Readings

# How easy is to represent preferences?

- Explicit representation of binary relations can be impossible in combinatorial domains.
- Logical representations are more compact, but only special cases allow to easy algorithmic solutions.
- Implicit representations (value functions) are easier to handle, but need much more careful and expensive modelling.

# How easy is to aggregate preferences?

- Only simple additive value functions are really easy to handle.
- Non linear value functions often require an exponential number of parameters to learn and may lead to non linear optimisation algorithms.
- Direct aggregation of binary relations through majority rules require careful hierarchical modelling in order to avoid combinatorial explosion.

# How easy is to learn preferences?

Generally speaking is complicated ...

- Direct protocols are cognitively complex.
- Indirect protocols are cognitively simple, but most of the times require complex (NP hard) algorithms or heuristics.
- Parameter estimation works fine only for linear models.

# How easy is to design learning protocols?

- Minimal number of questions.
- Cognitive burden for the client.
- Accuracy vs. constructive learning.

# Response Models

Model the user's cognitive ability of answering correctly to a preference query

- *Noiseless* responses (unrealistic but often assumed in research papers!)
- *Constant* error (can model distraction, e.g. clicking on the wrong icon)
- *Logistic* error (Boltzmann distribution), a commonly used probabilistic response model for comparison/choice queries: "*Among options in set S, which one do you prefer?*"

probability of response "x is my preferred item in S"

$$P_r(S \rightarrow x) = \frac{e^{\gamma u(x)}}{\sum_{y \in S} e^{\gamma u(y)}}$$

$\gamma$ is a temperature parameter (how "noisy" is the user).

For comparison queries ("Is item1 better than item2?")
P(selecting 1st item) as a function of the difference in utility

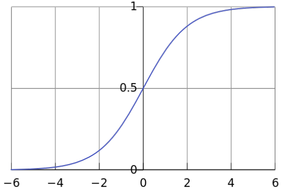

# What Query to Ask Next?

- The problem can be modeled as a POMDP *[Boutilier, AAAI 2002]*, however impractical to solve for non trivial cases
- Idea: ask query with highest "value", a posteriori improvement in decision quality
- In a Bayesian approach, (Myopic) *Expected Value of Information*

$$EVOI_\theta(q) = \sum_{r \in R} P_\theta(r) EU^*_{\theta|r} - EU^*_\theta$$

where $R$ is the set of possible responses (answers); $\theta$ is the current belief distribution and $\theta|r$ the posterior and $P_\theta(r)$ the prior probability of a given response.

- Ask query $q^* = \arg\max EVOI_\theta(q)$ with highest EVOI
- In non-Bayesian setting, one can use non probabilistic measures of decision improvement (for example, worst-case regret reduction, ...)

# Argumentation Theory

## Argumentation framework

An *argumentation framework AF* is a tuple $(Ar, \mathcal{R})$, where $Ar$ is a set of arguments and $\mathcal{R}$ is a binary relation on $Ar$ (i.e., $\mathcal{R} \subseteq Ar \times Ar$). An argument $A$ *attacks* an argument $B$ iff $(A, B) \in \mathcal{R}$.

## Conflict-free Extension

Let $(Ar, \mathcal{R})$ be an argumentation framework. The set $S \subseteq Ar$ is *conflict-free* if and only if there are no $A, B \in S$ such that $(A, B) \in \mathcal{R}$. We also say that $S$ *defends* $A$ (or, the argument $A$ is *acceptable with respect to* $S$) if, $\forall B \in Ar$ such that $(B, A) \in \mathcal{R}$, $\exists C \in S$ such that $(C, B) \in \mathcal{R}$.

# Semantics

## Acceptability semantics

Let $(Ar, \mathcal{R})$ be an argumentation framework and set $S \subseteq Ar$.

- $S$ is an *admissible* extension if and only if it is conflict-free and defends all its elements.
- $S$ is a *complete* extension if and only if it is conflict-free and contains precisely all the elements it defends, i.e., $S = \{A \mid S \text{ defends } A\}$.
- $S$ is a *grounded* extension if and only if $S$ is the smallest (w.r.t. set inclusion) complete extension of $AF$.
- $S$ is a *preferred* extension if and only if $S$ is maximal (w.r.t. set inclusion) among admissible extensions of $AF$.
- $S$ is a *stable* extension if and only if $S$ is conflict-free and $\forall B \notin S, \exists A \in S$ such that $(A, B) \in \mathcal{R}$.

# Why do we need that?

- Why $x \succeq y$?
- Why $x$ has been chosen?
- Why $y$ has not been chosen?

**Deep reasons**

Not only factual justifications, but also methodological ones.

# EURO WG Preference Handling

http://preferencehandling.free.fr/

Established in 2007 (but practically operating since 2004), coordinated by Ulrich Jünker, it organises an annual MPREF workshop (associated to major AI, DB or OR conferences) and maintains an active site full of information, tutorials, news etc..

EURO
The Association of European Operational Research Societies

# Algorithmic Decision Theory

http://www.algodec.org

- Formally initiated as a COST ACTION (http://www.cost-ic0602.org) started in 2007 and ended in 2011.
- Evolved as an international community addressing issues related to computational aspects of decision theory.
- Organised the $1^{st}$ (Venice, IT, 2009) and $2^{nd}$ (Rutgers, NJ, USA) International Conferences on *Algorithmic Decision Theory*.
- Organising $3^{rd}$ International Conference in Brussels, BE, the 13-15/11/2013 (http://www.adt2013.org).
- Proceedings available within the Springer-Verlag LNAI series (volumes 5783, 6992, 8176).

# Computational Social Choice

http://www.illc.uva.nl/COMSOC/

- Is an international community concerned by issues related to computational aspects of social choice theory and how it can be used in settings others than voting and collectively deciding.
- It is also a COST Action (http://www.illc.uva.nl/COST-IC1205/, started in 2012.
- Organising a series of International Workshops (Amsterdam, 2006;, Liverpool, 2008; Düsseldorf, 2010; Krakow, 2012).
- The next COMSOC International Workshop will take place at CMU, Pittsburgh the 23-25/06/2014 (http://www.cs.cmu.edu/∼arielpro/comsoc-14/).

# References 1 (Fundamentals: Preference Modeling)


- Roberts F.S, *Measurement theory, with applications to Decision Making, Utility and the Social Sciences*, Addison-Wesley, Boston, 1979.
- Roubens M., Vincke Ph., *Preference Modeling*, Springer Verlag, Berlin, 1985.
- Fishburn P.C., *Interval Orders and Interval Graphs*, J. Wiley, New York, 1985.
- Fodor J., Roubens M., *Fuzzy preference modelling and multicriteria decision support*, Kluwer Academic, Dordrecht, 1994.
- Pirlot M., Vincke Ph., *Semi Orders*, Kluwer Academic, Dordrecht, 1997.
- Fishburn P.C., "Preference structures and their numerical representations", *Theoretical Computer Science*, vol. 217, 359-383, 1999.
- Öztürk M., Tsoukiàs A., Vincke Ph., "Preference Modelling", in M. Ehrgott, S. Greco, J. Figueira (eds.), *State of the Art in Multiple Criteria Decision Analysis*, Springer Verlag, Berlin, 27 - 72, 2005.

# References 2 (Preference Aggregation and Learning)


- Arrow K.J., *Social choice and individual values"*, J. Wiley, 1951.
- Sen A.K., *Collective Choice and Social Welfare*, North Holland, 1970.
- Keeney R.L., Raiffa H., *Decisions with multiple objectives: Preferences and value tradeoffs*, J. Wiley, 1976.
- Arrow K.J., Raynaud H., *Social choice and multicriterion decision-making*, MIT Press, 1986.
- Wakker, P.P., *Additive Representations of Preferences: A new Foundation of Decision Analysis*, Kluwer Academic, 1989.
- Tanguiane, A.S., *Aggregation and Representation of Preferences*, Springer-Verlag, 1991.
- Kahneman D., Tversky A., *Choice, Values, Frames*, The Cambridge University Press, 2000.

# References 3 (More ...)


- von Wright G.H. *The logic of preference*, Edinburgh University Press, 1963.
- Pawlak Z., *Rough Sets - Theoretical Aspects of Reasoning about Data"*, Kluwer Academic, 1991.
- Boutilier C., Brafman R.I., Domshlak C., Hoos H.H., Poole D., "CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements", *Journal of Artificial Intelligence Research*, vol. 21, 135-191, 2004.
- Lang J., "Logical representation of preferences", in D. Bouyssou, D. Dubois, M. Pirlot and H. Prade (eds.), *Decision-Making Process: Concepts and Methods*, J. Wiley, 321-363, 2009.
- Janowitz M.F., *Ordinal and Relational Clustering*, World Scientific, 2010.
- Füurnkranz J., Hüllermeier E., *Preference Learning*, Springer Verlag, 2010.
- Modgil S., Prakken H., "A general account of argumentation with preferences' , *Artificial Intelligence*, vol. 195, 361-397, 2013.

# The Authors

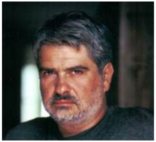

- Alexis Tsoukiàs is a research director at CNRS and the director of the *Laboratoire d'Analyse et Modélisation de Systèmes pour l'Aide à la DEcision* (LAMSADE) of Paris Dauphine University.
- Paolo Viappiani is a researcher at CNRS working at the computer science lab of the University Pierre et Marie Curie in Paris (UPMC).