Learning multiple multicriteria additive models from heterogeneous preferences

Vincent Auriau[A,B], Khaled Belahcène[A], Emmanuel Malherbe[B] and Vincent Mousseau[A]
[A]MICS, CentraleSupélec Université Paris-Saclay, Gif-sur-Yvette, France
[B]Artefact Research Center, Paris, France

Additive preference representation is standard in Multiple Criteria Decision Analysis, and learning such a preference model dates back from the UTA method [11]. In this seminal work, an additive piece-wise linear model is inferred from a learning set composed of pairwise comparisons. In this setting, the learning set is provided by a single Decision-Maker (DM), and an additive model is inferred to match the learning set. We extend this framework to the case where *(i)* multiple DMs with heterogeneous preferences provide part of the learning set, and *(ii)* the learning set is provided as a whole without knowing which DM expressed each pairwise comparison. Hence, the problem amounts to inferring a preference model for each DM and simultaneously "discovering" the segmentation of the learning set. In this paper, we show that this problem is computationally difficult. We propose a mathematical programming based resolution approach to solve this *Preference Learning and Segmentation* problem (PLS). We also propose a heuristic to deal with large datasets. We study the performance of both algorithms through experiments using synthetic and real data.

1. Introduction and motivation

Multiple Criteria Decision Analysis (MCDA) aims at supporting a Decision-Maker (DM) and at providing him/her recommendations concerning the decision tackled. These recommendations are based on a preference model that faithfully represents the DM's judgments. We consider, in this paper, an additive value-based representation of preferences [13]. For the additive model to faithfully represent the DM's preferences, adequate values should be assigned to the model parameters. This is usually done through a preference learning phase, which takes as input observed preferences such as pairwise comparisons [22].

A standard value-based setting for preference learning in multicriteria ranking problems with an additive representation of preferences dates back from the UTA method [11]. This seminal work considers a multicriteria ranking problem in which an additive piece-wise linear model is inferred from preferences. In this setting, a single DM provides the learning set, and the UTA method infers, with linear programming, the additive piece-wise linear model that best matches the learning set, even with noisy preferences.

Example 1 (Clients preferences in a supermarket): A supermarket is willing to adapt the list of products to its customer base. The products selected should align with the client's preferences. Yet, not all customers necessarily have the same preferences, and it is standard to consider a market segmentation in which each segment represents a group of clients with homogeneous preferences. Each store can rely on actual sales to identify the segmentation and learn the clients' preferences. For each product sale x, one can derive a comparison $x \succcurlyeq y$ for each product y sold in the store and substitutable to x. Hence, the problem can be formulated as clustering the clients and learning a preference model for each cluster.

Example 2 (Films comparisons on social media): A website about cinema presents several films and

Published at the International Conference on Algorithmic Decision Theory (ADT) 2024
Artefact Research Center, Paris, France

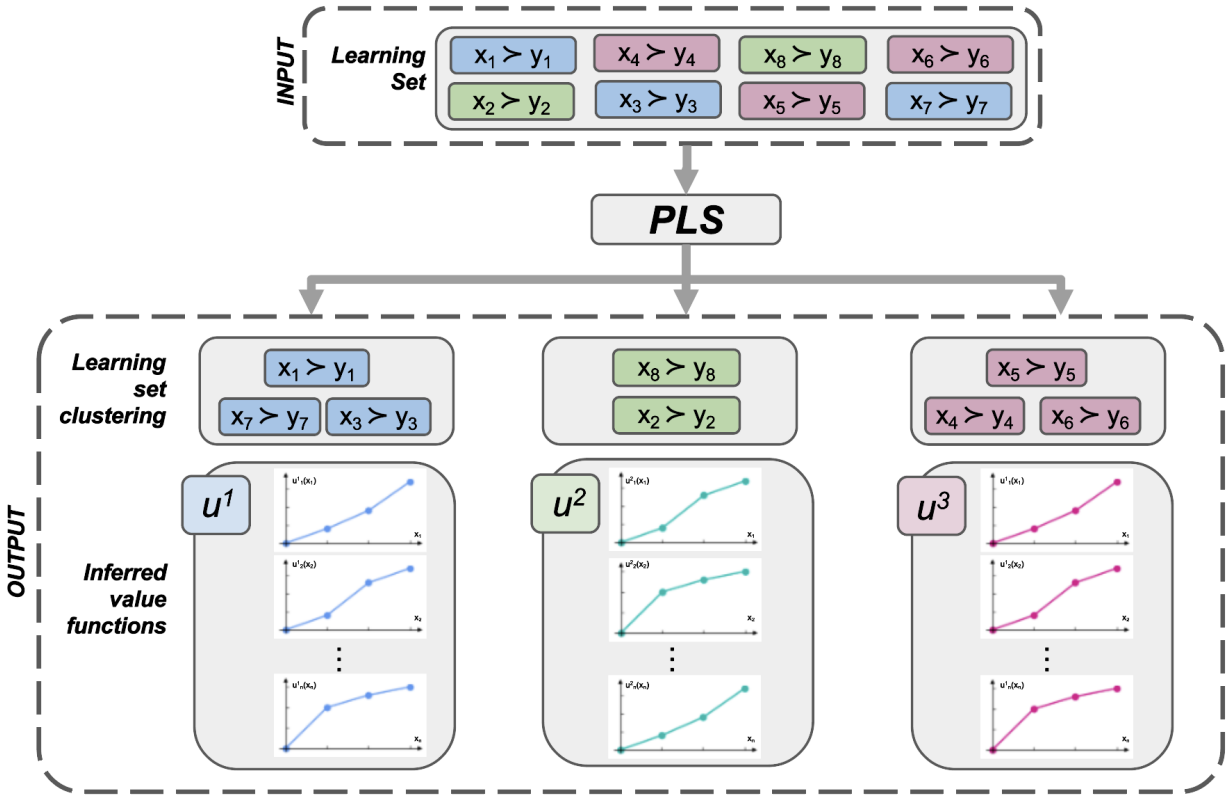


Figure 1 | Illustration of Preference Learning and Segmentation with 3 clusters

requests user feedback. Users provide film scores from which pairwise comparisons can be derived. Users are not requested to have an account and are hence unidentifiable. For marketing purposes, the website wants to identify, from the collected data, k clusters of users, with each a specific preference model.

Example 3 (Wine tasting): A group of oenologist experts is performing a blind tasting evaluation of various wines. These wines have been previously evaluated from a chemical point of view, and the concentration of several chemical components known to characterize the quality of the wines are measured. Each expert is required to attach to wines a number of "stars" (from ⋆ to ⋆⋆⋆⋆⋆). All experts in the panel obviously are not fully aligned in the way they assess wines. We would like to identify groups of experts with similar tastes and how each group accounts for wine characteristics in their preferences.

These examples correspond to instances of a *Preference Learning and Segmentation* (PLS) problem. Such PLS problem takes as input a learning set composed of a list of pairwise comparisons and should determine a given number of preference models so that each comparison in the learning set is correctly represented by at least one of the models. A key issue in PLS problems is that the partition of the learning set is not known and should be computed along the inferred preference models. Hence, the outputs of the PLS problem are both *(i)* the inferred preference models, and *(ii)* the corresponding partition of the learning set.

In this paper, we propose a formal setting for the PLS problem, usually not considered in preference learning for MCDA, in which decisions from heterogeneous sources are observed. We present two algorithms for preference learning segmentation: a MILO formulation and a heuristic method. Finally, we have conducted experiments to establish our model's performances using synthetic and real data.

The paper is organized as follows. In the next section, we review related works that are of interest with respect to our contribution. Section 3 is devoted to the formalization of the PLS problem, its resolution via a mathematical programming formulation, and the study of its computational complexity. A heuristic method specifically developed for the PLS problem is presented in Section 4. The experiments performed on synthetic data, in section 5, allows to analyze the performance of both proposed algorithms. In Section 6, we show the benefit of our contribution on a a-posteriori experiment based on real data. Final section groups conclusions and avenues for further research.

2. Related Work and Previous Research

2.1. MCDA and the UTA model

In multi-criteria decision analysis, alternatives are evaluated on several conflicting criteria. Formally, we consider a set $\mathcal{N} = \{1, \cdots, n\}$ of n criteria, where the evaluation scale on criterion i, $i \in \mathcal{N}$, is denoted $X_i \subset \mathbb{R}$. Hence an alternative $x = (x_1, ..., x_n)$ is an element of $X = \prod_{i \in \mathcal{N}} X_i$.

To compare such alternatives x and y, a value-based approach consists in defining a value function u such that $x \succcurlyeq y \Leftrightarrow u(x) \geq u(y)$, where $\succcurlyeq$ is a weak preference relation defined on X ($\succ$ and $\sim$ its symetric and asymetric part). In the following we consider w.l.o.g. that criteria are to be maximized, i.e., $x_i \succcurlyeq_i y_i \Leftrightarrow x_i \geq y_i$, where $\succcurlyeq_i$ denotes the weak preference relation on criterion i.

The UTA [11] method defines a value function $u(\cdot)$ using additive marginal value functions $u_i(\cdot)$, defined on X_i, $i \in \mathcal{N}$: $u(x) = \sum_{i \in \mathcal{N}} u_i(x_i)$, $x \in X$. On criterion $i \in \mathcal{N}$, the marginal value function $u_i(\cdot)$ is piecewise linear. More precisely, the scale of each criterion X_i is divided into L equal interval, $[x_i^0, x_i^1], [x_i^1, x_i^2], \ldots, [x_i^{L-1}, x_i^L]$, and the marginal value function $u_i(\cdot)$ is linear in each interval. Hence, $u_i(\cdot)$ is defined by its values at each breakpoint, denoted $u_i(x_i^l)$, $i \in \{1..n\}$, $l \in \{0..L\}$. The piecewise marginal functions $u_i(\cdot)$ when $x_i \in [x_i^l, x_i^{l+1}]$ can be written as:

$$u_i(x_i) = u_i(x_i^l) + \frac{x_i - x_i^l}{x_i^{l+1} - x_i^l} \cdot \left(u_i(x_i^{l+1}) - u_i(x_i^l)\right) \tag{1}$$

The parameters of the value model, $u_i(x_i^l) \in \mathbb{R}^+, i \in \mathcal{N}, l = 0..L$ are to be estimated using a learning set $\mathcal{P}$ provided by the DM. This learning set is composed of P pairwise comparisons:

$$\mathcal{P} = \{(x^{(1)}, y^{(1)}), (x^{(2)}, y^{(2)}), \cdots, (x^{(P)}, y^{(P)})\}$$

where $x^{(j)} = (x_1^{(j)}, ..., x_n^{(j)})$ is preferred to $y^{(j)} = (y_1^{(j)}, ..., y_n^{(j)})$, $j = 1..P$.

In such settings, the objective is to obtain the functions u_i so that:

$$x^{(j)} \succ y^{(j)} \Leftrightarrow \sum_i u_i(x_i^{(j)}) > \sum_i u_i(y_i^{(j)}) \tag{2}$$

The model follows additional normalization and monotonicity constraints:

$$u_i(x_i^l) \leq u_i(x_i^{l+1}), \forall i \in \mathcal{N}, \forall l = 0..L-1 \tag{3a}$$

$$u_i(x_i^0) = 0, \forall i \in \mathcal{N} \tag{3b}$$

$$\sum_{i=1}^{n} u_i(x_i^L) = 1 \tag{3c}$$

The model can be formalized as a Linear Optimization problem. It ensures an efficient and scalable computation using solvers. For this purpose, we introduce potential estimation errors, $\sigma^{(j)} \in \mathbb{R}^+$. They are used to add an additional preference constraint for all observed pairs $j \in \{1..P\}$:

$$u(x^{(j)}) - u(y^{(j)}) + \sigma^{(j)} \geq 0, \forall j \in \{1..P\} \tag{3d}$$

The optimization objective is formulated as the minimization of the estimation errors:

$$\min \sum_{j=1}^{P} \sigma^{(j)} \tag{UTA_{obj}}$$

Variants of the UTA model [24] exist for different types of data or optimality criteria. Along this adaptability and efficient computation, the method has been popularized by its interpretability potential.

2.2. Market Segmentation

Understanding its customer base is a critical stake for retailers [25]. It is used to adjust the marketing strategy as well as increasing satisfaction. A common approach is to delineate market segments: buyers with similar characteristics are grouped together and represent a segment of the customer base. The similarity between individuals is usually based on demographic, geographic or choice data. Natural tools to uncover these different segments are clustering algorithms such as k-means [28] or hierarchical clustering [16]. If using the customer characteristics is natural, as stated in [17], individual with similar socioeconomic attributes may have very different preferences. Introducing attitudinal information [17] helps taking into account the decision process of the customer. Such data can be collected through surveys which are difficult and expensive to carry out. Leveraging past transactions for modeling user choices has also been studied, in particular with the family of latent class models [27]. Recent work shows promising perspectives on large datasets [2], however this probabilistic modeling may not reveal distinct clusters of preferences. Engineering features [15] can also prove to be efficient but usually needs several data points for each customer so that it is possible to represent well its behavior.

In order to overcome these limitations, we propose to create segments based on customer decisions, taking into account the alternatives attributes. Specific information on the individuals is not needed although it could be integrated.

2.3. Clustering Rankings

When disagreements or incompatibilities are observed in the learning set, the UTA model is designed to find the best consensus. In order to integrate heterogeneity in the modelization, a preprocessing step can be carried out to partition the data into homogeneous sets. A standard method is to work with clustering algorithms based on the different decision makers attributes. However when few or no data is available or if decision makers have only been exposed to a subset of alternatives, an approach is to leverage their decision process with finer granularity. A common algorithm used with heterogeneous decision makers and ranking data is the collaborative filtering [8]. The algorithm is able to generalize from partial evaluations of alternatives by individuals. It leverages a proximity distance between individuals based on their common evaluations of the same alternatives. Several methods have proven their efficiency to create segments of the population based on a collaborative filtering output [26]. Another possibility is to define a similarity measure between two rankings that

can be used to create clusters. [4] leverages the Mahalanobis distance to create groups of decisions within a defined dissensus threshold.

These methods do not describe the different alternatives on relevant criteria. An advantage is that the data collection process is simpler. However one might want to generalize a model to unknown alternatives or to take into account criteria that may vary over time such as the price. For this purpose, [5, 7] propose to compute a ranking function by decision maker and then to clusterize these value functions. In order to use clustering algorithms, it is needed to define the appropriate distance measure in the value function space. [7] uses such approach for a search engine that generates recommendations based on a click history and [5] regroups spectators according to their movie tastes.

These methods suppose that enough preferences are collected for every user. In the context of a webpage or a supermarket, it is difficult to track customers. The database may be an agglomeration of anonymous contributions, that must therefore be handled differently. Latent class models have been developed for such cases. For instance, mixtures of Plackett-Luce models [9, 19] or Mallows [3, 18] have been employed to model elections. It is a typical case of an unidentified population that has extremely heteregenous tastes. However, these mixture models consider defined and fixed alternatives. For instance, candidates are well identified, the clusters are not based on the different elements of their programs but the rankings of the candidate themselves. Going back to our introductive *Example 3*, it would mean that the computed model would not be able to generalize to an untasted wine or to the next year when wines will be different. The experts decision process is not learned nor used. Furthermore some criteria such as the price can prove to be particularly important in the decision making of customers. In *Example 1* it has a direct impact on customers preferences and can change. It is therefore compulsory to integrate such information. Moreover, obtaining parameters such as a sensitivity toward a criterion will not only better match the customer decisions but also provide meaningful insight.

Finally, [20] uses a hierarchical model to represent heteogenerous consumers that is estimated with Bayesian inference. Priors are defined and adjusted to represent each customer membership to a customer segment and then its preferences. While this model is able to work with large datasets, it mainly relies on the choice of the right priors. This prior selection is mainly subjective and lacks an efficient validation method.

3. Preference Learning and Segmentation (PLS) problem

3.1. Problem setting

We consider a heterogeneous set of P observed preferences. These preferences are expressed as a pair of alternatives $(x^{(j)}, y^{(j)})$, $j \in \{1..P\}$, where $x^{(j)}$ has been preferred to $y^{(j)}$ by an unknown decision maker. The alternatives are described on n criteria whose evaluation scale is noted X_i for criterion $i \in \{1..n\}$.

The Preference Learning and Segmentation problem aims to cluster the P pairs into K clusters and to learn a multicriteria preference model for each of these clusters. A particular case occurs when these preference models are represented by an additive value function, i.e., when we learn a value function u^k on each cluster. An obvious intent of the PLS is to represent the maximum number observed preferences. By that, we mean that for all $(x^{(j)}, y^{(j)})$, $j \in \{1..P\}$, at least one of the k value functions expresses the observed decisions or differently written:

$$x^{(j)} \succcurlyeq^K y^{(j)} \iff \exists k \in \{1..K\} \text{ such that } \sum_i u_i^k(x_i^{(j)}) \geq \sum_i u_i^k(y_i^{(j)}) \tag{4}$$

3.2. Mixed Integer Linear Optimization resolution approach

We propose a Mixed Integer Linear Optimization (MILO) approach to solve our Preference Learning and Segmentation problem. The main objective is to cover at best the observed preferences with K different UTA models. In the following, the models are noted $u^k(.), k \in \{1..K\}$ and their marginal values on criterion $i, u_i^k(.)$.

As similarly introduced in Section 2.1, the $u^k, k \in \{1..K\}$ value functions of the UTA model are represented by their marginal coefficients on criterion $i \in \{1..n\}$, $u_i^k(x_i^l), l \in \{1..L\}$, to be computed. They also follow the same normalization and monotonicity constraints defined in Equations 3a, 3b and 3c $\forall k \in \{1..K\}$:

$$\sum_{i=1}^{n} u_i^k(x_i^L) = 1 \tag{5a}$$

$$u_i^k(x_i^0) = 0, \forall i \in \{1..n\} \tag{5b}$$

$$u_i^k(x_i^l) \leq u_i^k(x_i^{l+1}), \forall i \in \{1..n\}, \forall l \in \{0..L-1\} \tag{5c}$$

Additionally to compute the marginal coefficients, our optimization formulation attributes each observed preference to one of the K models. For this purpose, we introduce the binary variables $z^{k,(j)}, j \in \{1..P\}$ and $k \in \{1..K\}$ as well as the following constraints $\forall j \in \{1..P\}$:

$$\sum_{k=1}^{K} z^{k,(j)} \geq 1 \tag{6a}$$

$$\sum_{i=1}^{n} u_i^k(x_i^{(j)}) - \sum_{i=1}^{n} u_i^k(y_i^{(j)}) + M \cdot (1 - z^{k,(j)}) + \sigma^{(j)} \geq 0, \forall k \in \{1..K\} \tag{6b}$$

with M a majorant of sufficient magnitude and $\sigma^{(j)} \geq 0$ estimation errors. The constraint in Equation 6a makes sure that every pair is attributed to a cluster. The constraint in Equation 6b ensures that when a pair is attributed to cluster k, the corresponding value function u^k respects the preference. The linear interpolation of $u_i^k(x_i)$ between $u_i^k(x_i^l)$ and $u_i^k(x_i^{l'})$ is defined as in (1).

Finally we formulate the following optimization objective, minimizing the total sum of estimation errors:

$$\min \sum_{j=1}^{P} \sigma^{(j)} \tag{PLS_{obj}}$$

This formulated MILO can be solved using optimization solvers. In terms of complexity, it handles $K \times P$ binary variables as well as $K \times n \times L$ continuous variables in $[0; 1]$. The complexity grows with the number of clusters to be estimated, but most of all with the number of observed preferences. In the experiments, we provide solving time insights on different instance sizes.

The presented model defines a set of binary variables $z^{k,(j)}$ for each of the observed pairwise comparisons. One can ensure that groups of pairs are clusterized together using the same binary variables in 6b. For instance, as further described in Section 5.2, a choice can be broken into several pairwise comparisons to be clusterized. The $z^{k,(j)}$ binary variables are set to be equal for these pairs.

3.3. Problem difficulty

In this section, we study the problem of rationalizing a set of observed preferences by segmenting them and finding a model for each segment from a computational perspective. When the number of segments

is one and the sought model is the additive value model with piecewise linear marginal utilities with fixed breakpoints, the UTA method reformulates an instance of the 1-DM UTA SEGMENTATION problem as an instance of LINEAR OPTIMIZATION (see Section 2.1) which can be solved in polynomial time [23]. When allowing for several DMs, the problem involves partitioning the data and seems more arduous. We indeed show that, even in the simplest configuration where each criterion is evaluated on a binary scale, the problem is NP-complete for any fixed value of $K \geq 2$.

We begin by giving a proper definition of this trimmed down version, where the observed preferences can be interpreted as subset comparisons. We name the problem by referring to the majority rule because when trying to represent preferences with an additive value function, each criterion is given a weight that can be interpreted as a voting power.

MULTIPLE DECISION-MAKER MAJORITY RULE SEGMENTATION (K-DM MRS)	
Input:	set N of criteria, set $\mathcal{P}$ of comparisons between subsets of N, positive integer $K < \lvert N \rvert$
Question:	is there a partition of $\mathcal{P}$ into K disjoint subsets $\mathcal{P}^1, \dots, \mathcal{P}^K$ and K functions $W_1, \dots W_K : N \to \mathbb{Q}^+$ such that, for each statement $A \succ B$ in $\mathcal{P}^k$ we have $\sum_{a \in A} W_k(a) > \sum_{b \in B} W_k(b)$?

Membership of this problem to the class NP is obvious, as the partition and weights form a polynomial certificate by design. We detail the hardness results, differentiating between the case with 3 or more decision-makers, where we directly interpret the K-partition of preference statements as a K-coloring, from the case with 2 decision-makers, as obtaining a 2-coloring is polytime but we leverage the 2-partition of preference statements to mimic conditional statements allowing to represent the constraints of a vertex cover problem. Complete proofs can be found in the supplementary material.

Proposition 3.1. *for $K \geq 3$, K-DM MRS is NP-complete.*

sketch. reduction from K-COLORABILITY [6, 12]. Given a graph $\mathcal{G} = (V, E)$, we produce an instance $(N, \mathcal{P})$ of K-DM MRS as follows.

- *Criteria:* let $N := V \cup \{\tau_0, \tau_1, \tau_2\} \cup \bigcup_{k \in [K]} \{\omega_k\}$. We note $\Omega_{-k} := \bigcup_{k' \in [K] : k' \neq k} \{\omega_k\}$ and $N_{-k} := N \setminus \{\omega_k\}$.
- *Preference statements:* let $\mathcal{P} := \bigcup_{v \in V} \pi_v \cup \bigcup_{k \in [K]} \mathcal{P}_k^E \cup \bigcup_{k \in [K]} \mathcal{P}_k^\tau \cup \bigcup_{k \in K} \pi_k$ where:
  - for all $v \in V$, π_v denotes the statement $\{v\} \succ \{\tau_1\}$;
  - for all $k \in [K]$, $\mathcal{P}_k^E$ and $\mathcal{P}_k^\tau$ denote the following collections of statements $\mathcal{P}_k^E := \bigcup_{\{u,v\} \in E} \{\{\tau_2\} \succ \{u, v\} \cup \Omega_{-k}\}$
    $\mathcal{P}_k^\tau := \{\{\tau_1\} \succ \{\tau_0\} \cup \Omega_{-k}, \{\tau_0, \tau_1\} \succ \{\tau_2\} \cup \Omega_{-k}\}$
  - finally, for all $k \in [K]$, π_k^ω denotes the statement $\{\omega_k\} \succ N_{-k}$.
- *Claim:* $\mathcal{G}$ is a positive instance of K-COLORABILITY $\iff$ $N, \mathcal{P}$ is a positive instance of K-DM MRS.
- *Hint:* observe the statements $\bigcup_{k \in [K]} \pi_k^\omega$ are pairwise inconsistent, so each segment contains exactly one of those, defining the color k associated to the vertices v whose statement π_v appear in the same bundle as π_k^ω. Statements in $\mathcal{P}_k^E$ or $\mathcal{P}_k^\tau$ are pairwise inconsistent with $\pi_{k'}^\omega$ with $k \neq k'$ thus appear in the same segment as π_k^ω and ensure the coloring is correct.

□

Proposition 3.2. *2-DM MRS is NP-complete.*

sketch. reduction from VERTEX COVER [6, 12]. Given a graph $\mathcal{G} = (V, E)$ and a positive integer $K < |V|$, we produce an instance $(N, \mathcal{P})$ of 2-DM MRS as follows.

- *Criteria:* Let $N := V \cup \{\alpha, \beta, \omega_1, \omega_2\} \cup \{\gamma_1, \ldots, \gamma_K\}$, such that $|N| = |V| + K + 4 < 2|V| + 4$.
- *Preference statements:* let $\mathcal{P} := \bigcup_{v \in V} \{\pi_v^+, \pi_v^-\} \cup \bigcup_{\{u,v\} \in E} \pi^E_{\{u,v\}} \cup \{\pi_2^\alpha, \pi_1^\omega, \pi_2^\omega\} \cup \mathcal{P}_1^\alpha$ where:
  - for all vertices $v \in V$, $\pi_v^+ : \{v\} \succ \{\alpha\}$ and $\pi_v^- : \{\beta\} \succ \{v\}$;
  - for all edges $\{u, v\} \in E$, $\pi^E_{u,v} : \{u, v\} \succ \{\alpha, \beta, \omega_2\}$;
  - $\pi_1^\omega : \{\omega_1\} \succ V \cup \{\omega_2, \alpha, \beta\}$, $\pi_2^\omega : \{\omega_2\} \succ V \cup \{\omega_1, \alpha, \beta\}$ and $\pi_2^\alpha : \{\alpha\} \succ \{\beta, \omega_1\}$;
  - $\mathcal{P}_1^\alpha : \Big\{\{\alpha\} \succ \{\beta, \omega_2\}, \{\alpha\} \cup \bigcup_{k \in [K]} \{\gamma_k\} \succ V \cup \{\omega_2\}\Big\} \cup \bigcup_{k=1}^K \{\{\alpha\} \succ \{\gamma_k, \omega_2\}\}$.
- *Claim:* $\mathcal{G}$ has a VERTEX COVER (i.e. a subset $V' \subset V$ such that for each edge $\{u, v\} \in E$ at least one of u and v belongs to V') of size at most $K \iff N, \mathcal{P}$ is a positive instance of 2-DM MRS.
- *Hint:* Observe the statements π_1^ω and π_2^ω cannot appear in the same segment. Thus, comparative statements with ω_1 (resp. ω_2) at the RHS are bundled together. Meanwhile, for each vertex $v \in V$, statements π_v^+ and π_v^- are pairwise inconsistent and appear in distinct bundles. The set $\{v \in V : \pi_v^+$ appears in the same segment as $\pi_1^\omega\}$ is a vertex cover of size at most K.

□

Hence, finding an exact solution to the PLS problem is difficult ; the MILO formulation presented in section 3.2 is therefore expected to be relevant for instances up to a certain size. The detailed proofs are shared in the online appendices[1].

4. Heuristic

In order to mitigate the difficulty to scale up the MILO formulation, we have developed a first heuristic that can work with large number of observed preferences. Similar to the K-Means algorithm [14], it consists of a succession of two steps: assignment and update.

Initialization:
K - value chosen by the user - clusters are initialized with the KMeans++ [1] procedure:

- Choose randomly one of the data points as the first cluster center
- Compute the Euclidian distance $D(x^{(j)}, y^{(j)}) = ||(x^{(j)} - y^{(j)}) - (x^* - y^*)||_2$ between each data point $(x^{(j)}, y^{(j)})$ and the closest cluster center (x^*, y^*)
- Choose randomly the next cluster center among the remaining data points with a probability proportional to $D(x)^2$
- Once K cluster centers are chosen, assign each data point to the cluster whose center is the closest

After the initialization, we proceed to an update step and then alternate assignment and update steps until convergence.

Assignment Step:
The assignment step attributes a cluster to each of the data point $(x^{(j)}, y^{(j)})$. Considering K UTA models computed in the update step, the data point is associated to the cluster $k = \arg\max_{k \in \{1..K\}} \left(u^k(x^{(j)}) - u^k(y^{(j)})\right)$.

Update step: Given a partitioning in K cluster of our learning set, we proceed to compute a UTA model u^k, as defined in Section 2.1, on each of the partitions.

Convergence: We consider that the algorithm has converged when the partition does not change between two assignment steps. A softer stopping criterion can be set through a maximum number of

[1] Appendices available at github.com/artefactory/learning-heterogeneous-preferences

successive iterations with the number of data points attributed to different cluster under a negligible ratio of the learning set.

The heuristic consists on a succession of Linear Programming problems solving, ensuring better scalability than our MILO formulation.

5. Empirical assessment of algorithms

We present experimental results of our two algorithms on synthetic data. The different algorithms have been implemented in Python. We have used the Gurobi solver [10] to solve the different optimization problems. The codebase is shared at github.com/artefactory/learning-heterogeneous-preferences.

5.1. Experiments on synthetic datasets

5.1.1. Experimental objectives

Synthetic data can be leveraged to better understand a model's behavior and performance. In particular, the synthetic experiments are designed toward three objectives:

- Assessing descriptive power of the algorithm, i.e., its ability to learn from heterogeneous preferences compared to a "single preference" formulation such as *UTA*.
- Assessing predictive power of the algorithm, i.e., its ability to generalize to unseen data.
- Assessing the scalability of the model with large datasets.

5.1.2. Experimental design

The following protocol has been carried out for the synthetic experiments:

Data generation

- Select the number of ground truth clusters, $K \in \{2, 3, 4\}$, the number of criteria, $n \in \{4, 8, 10\}$, the number of pairs in the learning set $P \in \{128, 256, 512, 1024, 2048, 4096\}$, and the noise level $NL \in \{0, 0.05, 0.1, 0.2\}$,
- Draw uniformly at random K additive piecewise linear functions following *UTA*'s normalization constraints,
- Draw uniformly at random pairs $(x^{(j)}, y^{(j)}) \in \prod_{i \in \mathcal{N}} X_i \times \prod_{i \in \mathcal{N}} X_i$ ($\frac{P}{K}$ for learning and 8192 for testing),
- with probability $1 - NL$, place the correct preference in the dataset (i.e., $a \succ b$ when $u^k(a) > u^k(b)$ for any pair (a, b)), and the opposite preference with probability NL.

Model Computation Knowing the right number of clusters, each instance is solved using the MILO formulation as well as using the heuristic.

Evaluation

For each instance, we evaluate the computed models on the learning set and test set. The *ExplainedPairs* metric estimates the ratio of pairwise comparisons that are correctly restored by at

least one the models.

$$ExplainedPairs = \frac{1}{P}\sum_{j=1}^{P} \mathbb{1}\{\exists k \leq K \mid u^k(x^{(j)}) \geq u^k(y^{(j)})\}$$

The full procedure is repeated 100 times for each combination of parameters.

5.1.3. Results and discussion

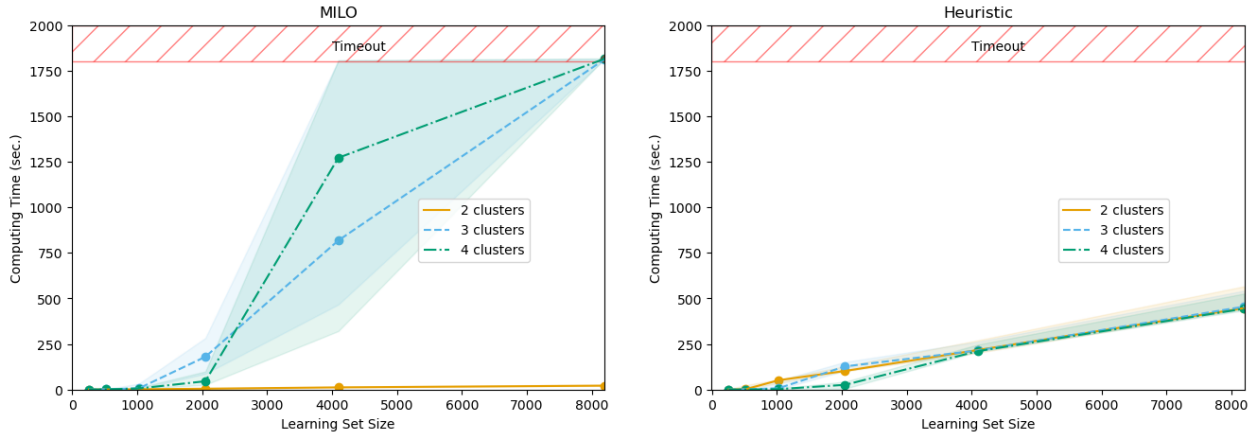


Figure 2 | Computing time for 8 criteria, 0% noise and different number of clusters. Shaded areas represent the values of first and 3rd quartiles over experiment repetitions.

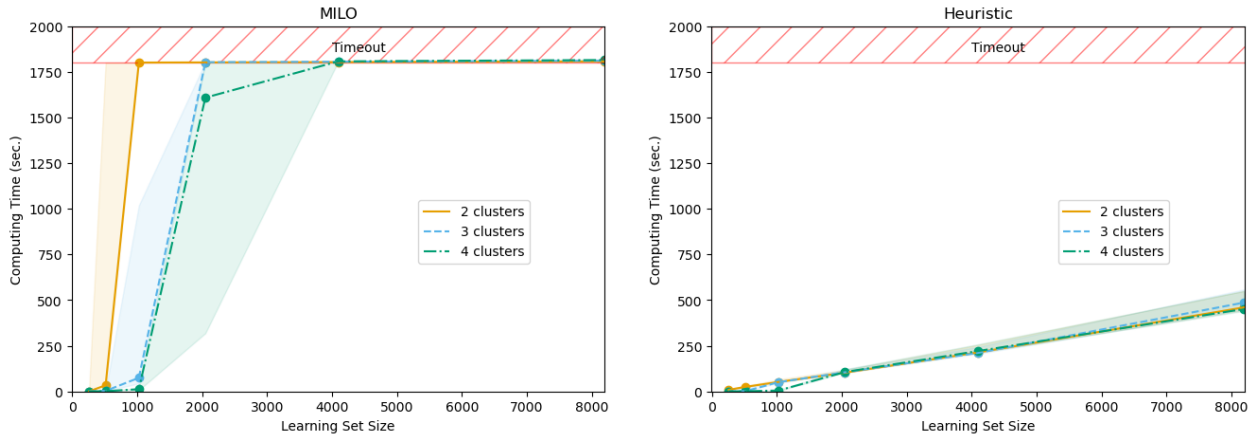


Figure 3 | Computing time for 8 criteria, 5% noise and different number of clusters. Shaded areas represent the first and 3rd quartiles over experiment repetitions.

In Fig. 2, we look into the computing time of our two models for different learning set sizes. We observe that the MILO formulation, quickly takes significant amounts of time, particularly when noise is added to the dataset, in Fig. 3. On the contrary the heuristic solving time seems to only increase linearly with the data quantity and seems not to be impacted by the chosen number of clusters.

We explore the ability of our MILO and heuristic to learn and generalize the heterogeneous preferences with the *ExplainedPairs* metric computed on the test Set. In Fig. 4, we observe the variations with different numbers of clusters and learning set sizes and in Fig. 5, the variations with

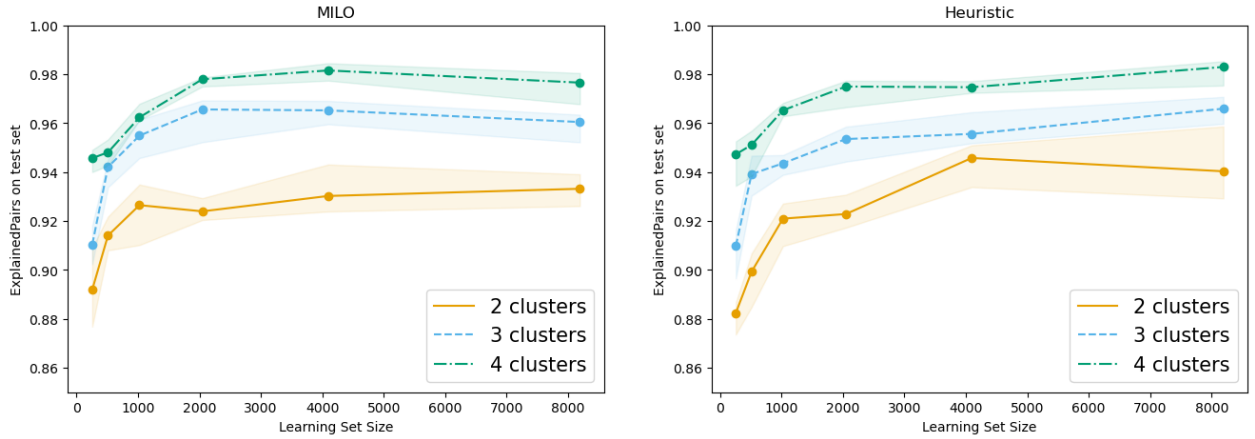


Figure 4 | Proportion of ExplainedPairs (%) on test set with 8 criteria, 5% noise and different cluster numbers. Shaded areas represent the first and 3rd quartiles over experiment repetitions.

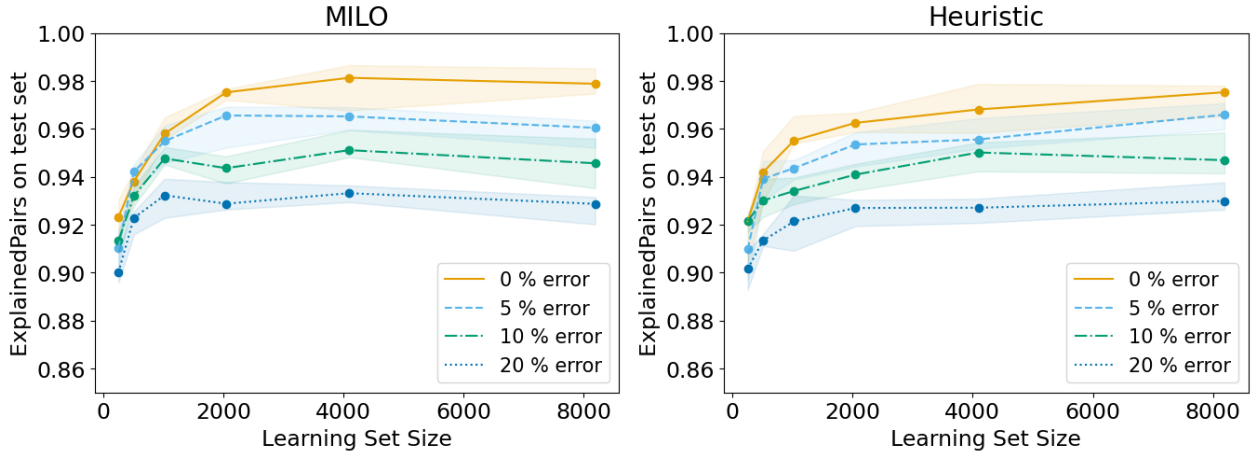


Figure 5 | Proportion of ExplainedPairs (%) on test set for 3 clusters, 8 criteria and different noise levels. Shaded areas represent the first and 3rd quartiles over experiment repetitions.

different noise levels. While we can explain some performance drops with the fact that the MILO reaches the defined timeout, we can globally observe that the more data, the better. In a similar expected fashion, the less noise in the learning set, the better the performances are. As a general comment, we consider that our approaches fullfill their main task, namely learning heterogeneous preferences. Additional results and settings are shared in the online appendices [2].

5.2. Experiments on real datasets

We investigate the performances of our algorithms on the car preferences dataset [3] [21]. This dataset results from a study conducted in California where respondents were asked to choose a car among six alternatives. The cars are evaluated on eight monotonous criteria described in Table 1. The feasibility of our approach is not guaranteed, in particular, the appropriate number of clusters needs to be

[2]Appendices available at github.com/artefactory/learning-heterogeneous-preferences
[3]Dataset available at github.com/artefactory/choice-learn

questioned.

5.2.1. Data processing

Different processing are applied to the dataset before estimating the model.

- alternatives with out-of-distribution criteria values are removed,
- preference pairs with dominance are removed from the dataset,
- criteria are normalized in $[0, 1]$ with 1 being the best value possible for each criterion.

The original dataset that includes 4654 choices, each involving 6 alternatives, leads to 18.498 pairwise comparisons.

5.2.2. Use case adaptation

The original dataset contains a set of choices made on sets of 6 alternatives. We exploit these choices as follows. The choice of x^1 in the set $\{x^1, x^2, \ldots, x^6\}$ results in the 5 following pairwise comparisons: $x^1 \succ x^2$, $x^1 \succ x^3$, $x^1 \succ x^4$, $x^1 \succ x^5$, and $x^1 \succ x^6$. Conversely, a computed model succeeds at restoring the choice of x^1 among $\{x^1, x^2, \ldots, x^6\}$ iff the 5 corresponding comparisons are simultaneously restored.

5.2.3. Experimentations

The dataset is randomly split in two: 10.000 randomly selected pairwise comparisons constitute a basis from which a number of learning sets of varying sizes are sampled. The remaining comparisons compose the test set. The comparisons derived from the same choice are grouped together. Both approaches are computed with $K \in \{1, 2, 3, 4, 5\}$ clusters on each of the learning sets. The experimental design is repeated 5 times with different random seeds.

5.2.4. Results

First, the resulting computing times lead to similar conclusions with regards to the synthetic experiments. The MILO formulation needs more than 6 hours with a few hundreds of pairs while the heuristic finds satisfying results within an hour for the largest set size, 10.000 pairs.

Fig. 6 showcases the heuristic performances on the test set with different learning set sizes. The performances are analyzed at the pair level, on the left, or at the choice level, on the right. First, we can observe that a single UTA model (1 cluster) cannot restore more than 64% of the preferences and 48% of the choices. The heterogeneity of the data cannot be well handled, showcasing the need for

Criterion	Description	Objective
price	Price of vehicle divided by the logarithm of income	min
range	Hundreds of miles vehicle can travel between refuelings/rechargings	max
acc	Acceleration, tens of seconds required to reach 30 mph from stop	min
speed	Highest attainable speed in hundreds of mph	max
pollution	Tailpipe emissions as fraction of those for new gas vehicle	min
space	Fraction of luggage space in comparable new gas vehicle	max
cost	Cost per mile of travel (tens of cents)	min
station	Fraction of stations that can refuel/recharge vehicle	max

Table 1 | Car preferences criteria description.

preference segmentation. A recurring problem in clustering is finding the right number of clusters representing the data. While several methods exist, such as using the silhouette score or the elbow method, it usually comes to a human-informed choice from the experimental observations. From Fig. 6, we can infer that 2 clusters are enough to clusterize the pairwise comparisons. However, at the choice level, 4 clusters seem like a better choice, with 90% of representation of the choices from the test set. The improvement brought by the fifth cluster is too marginal to be considered as the right value.

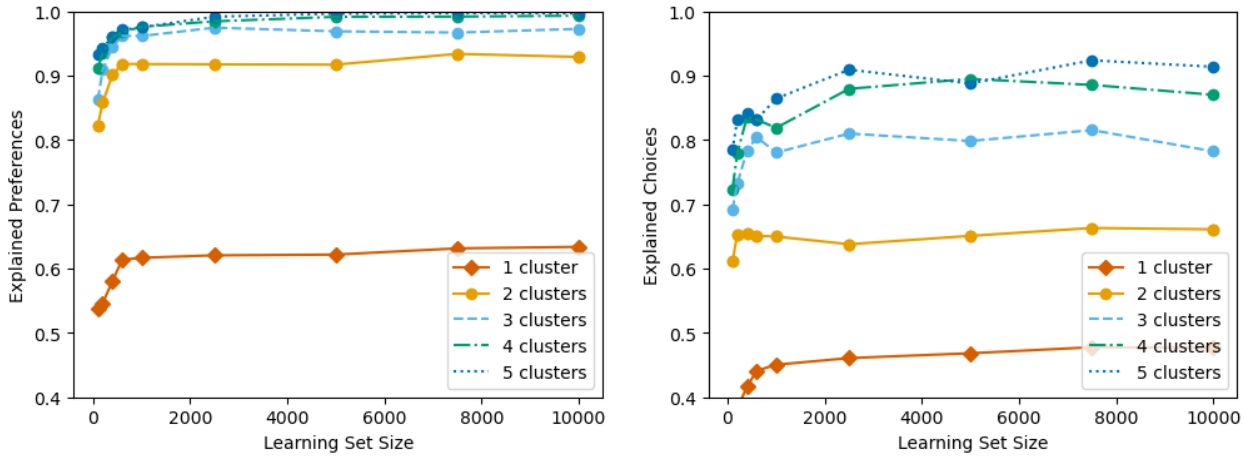


Figure 6 | Average *ExplainedPairs* and *ExplainedChoices* for the heuristic and different settings.

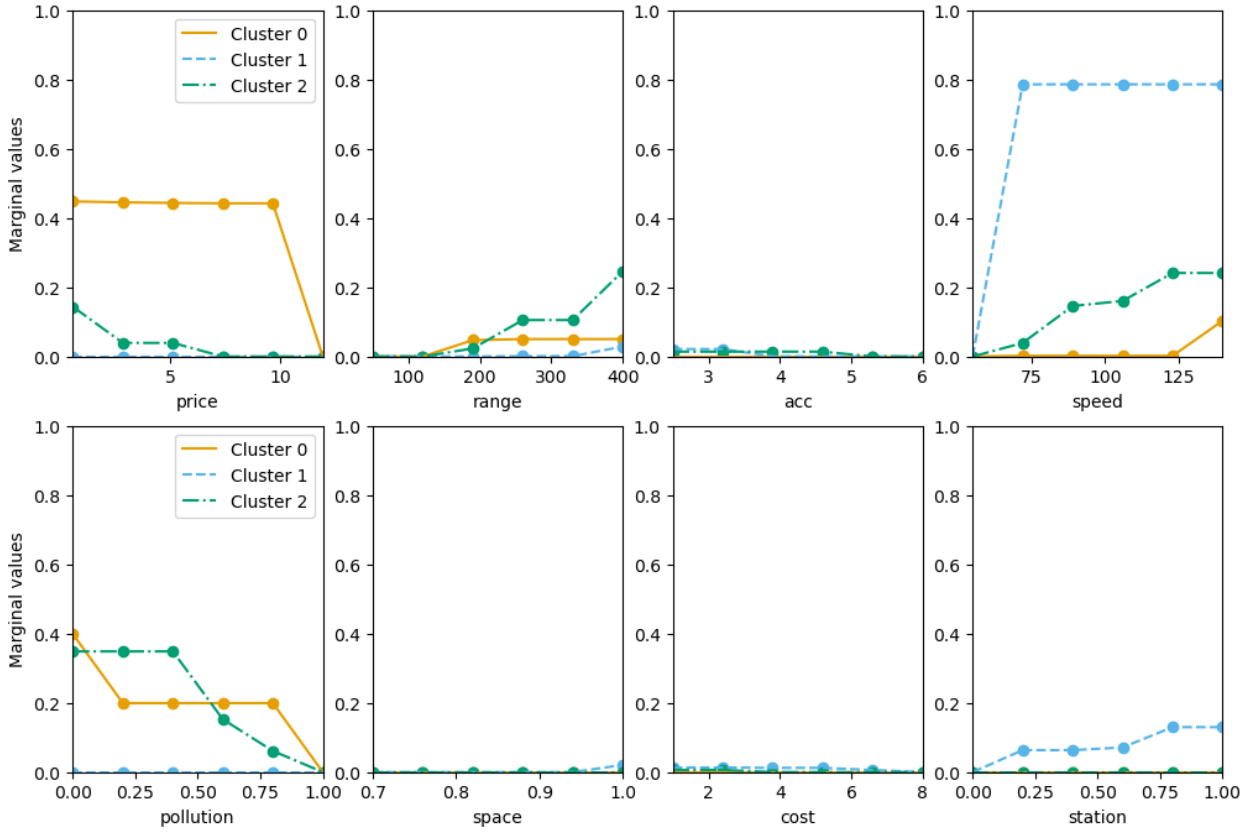


Figure 7 | Example of resulting UTA marginal values on the stated car preferences dataset.

Fig. 7 provides an example of resulting model coefficients for $K = 3$ clusters. A salient strength of our formulation is the direct interpretability of the compute models. We can observe how each criterion weighs on the decision with the coefficient values. For instance, the cluster 1 uses the price and pollution criteria to take a decision, while the cluster 2 mainly uses the speed criterion. Finally, the cluster 3, uses 4 criteria, with similar importance to represent the preferences.

6. Conclusion

In this paper we propose and study the *Preference Learning and Segmentation* (PLS) problem based on pairwise comparisons. In this problem, we want to clusterize the observed preferences and to compute a value function on each of the clusters. We show that the PLS problem is computationally difficult. We propose a mathematical programming as well as a heuristic resolution approach. We present experimental results on synthetic and real-world data.

These first results give rise to future research directions. First, we would like to use our approach in actual applications such as customer segmentation in supermarkets. Second, we plan to extend our approach to be able to specify PLS algorithms when replacing the additive preference model with alternative preference representations, such as pairwise comparison methods. Another interesting extension corresponds to a situation in which observed preferences that we consider as input are not comparisons but assignment examples, in which case we should move from choice problem formulation to sorting problems.

References

[1] Arthur, D., Vassilvitskii, S., et al.: k-means++: The advantages of careful seeding. In: Soda. vol. 7, pp. 1027–1035 (2007)

[2] Auriau, V., Aouad, A., Désir, A., Malherbe, E.: Choice-learn: Large-scale choice modeling for operational contexts through the lens of machine learning. Journal of Open Source Software (2024)

[3] Busse, L.M., Orbanz, P., Buhmann, J.M.: Cluster analysis of heterogeneous rank data. In: Proc. of the 24th internat. conf. on Machine learning. pp. 113–120 (2007)

[4] Cascon, J., González-Arteaga, T., de Andres Calle, R.: A new preference classification approach: The λ-dissensus cluster algorithm. Omega **111**, 102663 (2022)

[5] Díez, J., del Coz, J.J., Luaces, O., Bahamonde, A.: Clustering people according to their preference criteria. Expert Systems with Applications **34**(2), 1274–1284 (2008)

[6] Garey, M.R., Johnson, D.S.: Computers and intractability, vol. 174. Freeman San Francisco (1979)

[7] Giannopoulos, G., Brefeld, U., Dalamagas, T., Sellis, T.: Learning to rank user intent. In: Proc. of the 20th ACM international conference on Information and knowledge management. pp. 195–200 (2011)

[8] Goldberg, D., Nichols, D., Oki, B.M., Terry, D.: Using collaborative filtering to weave an information tapestry. Commun. ACM **35**(12), 61–70 (1992)

[9] Gormley, I.C., Murphy, T.B.: Exploring voting blocs within the irish electorate: A mixture modeling approach. Journal of the American Statistical Assoc. **103**(483), 1014–1027 (2008)

[10] Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2023)

[11] Jacquet-Lagreze, E., Siskos, J.: Assessing a set of additive utility functions for multicriteria decision-making, the UTA method. EJOR **10**(2), 151–164 (1982)

[12] Karp, R.: Reducibility among combinatorial problems. Complexity of Computer Computations pp. 85–104 (1972)

[13] Keeney, R.L., Raiffa, H.: Decisions with multiple objectives: preferences and value trade-offs. Cambridge University Press (1993)

[14] Krishna, K., Murty, M.N.: Genetic k-means algorithm. IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics) **29**(3), 433–439 (1999)

[15] Lefait, G., Kechadi, T.: Customer segmentation architecture based on clustering techniques. In: 2010 Fourth International Conference on Digital Society. pp. 243–248 (2010)

[16] Li, J., Wang, K., Xu, L.: Chameleon based on clustering feature tree and its application in customer segmentation. Annals of Operations Research **168**(1), 225–245 (2009)

[17] Li, Z., Wang, W., Yang, C., Ragland, D.R.: Bicycle commuting market analysis using attitudinal market segmentation approach. Transportation Research Part A: Policy and Practice **47**, 56–68 (2013)

[18] Liu, A., Moitra, A.: Efficiently learning mixtures of Mallows models. In: 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS). pp. 627–638 (2018)

[19] Liu, A., Zhao, Z., Liao, C., Lu, P., Xia, L.: Learning Plackett-Luce mixtures from partial preferences. Proc. of the AAAI Conf. on Artificial Intelligence **33**(01), 4328–4335 (2019)

[20] Liu, J., Wang, Y., Kadziński, M., Mao, X., Rao, Y.: A multiple criteria bayesian hierarchical model for analyzing heterogeneous consumer preferences. Omega p. 103113 (2024)

[21] McFadden, D., Train, K.: Mixed MNL models for discrete response. Journal of Applied Econometrics **15**(5), 447–470 (2000)

[22] Mousseau, V., Pirlot, M.: Preference elicitation and learning. EURO Journal on Decision Processes **3**(1-3) (2015)

[23] Nesterov, Y., Nemirovskii, A.: Interior-point polynomial algorithms in convex programming. SIAM (1994)

[24] Siskos, Y., Grigoroudis, E., Matsatsinis, N.F.: UTA methods. Multiple criteria decision analysis: State of the art surveys pp. 315–362 (2016)

[25] Tynan, A.C., Drayton, J.: Market segmentation. Journal of Marketing Management **2**(3), 301–335 (1987)

[26] Ungar, L.H., Foster, D.P.: Clustering methods for collaborative filtering. In: AAAI workshop on recommendation systems. vol. 1, pp. 114–129. Menlo Park, CA (1998)

[27] Wang, Q., Yang, X., Song, P., Sia, C.L.: Consumer segmentation analysis of multichannel and multistage consumption: A latent class mnl approach. Journal of Electronic Commerce Research **15**(4), 339 (2014)

[28] Zakrzewska, D., Murlewski, J.: Clustering algorithms for bank customer segmentation. 5th International Conferebce on Intelligent Systems Design and Applications pp. 197–202 (2005)