



HAL
open science

Cautious Learning of Multiattribute Preferences

Hugo Gilbert, Mohamed Ouaguenouni, Meltem Ozturk, Olivier Spanjaard

► **To cite this version:**

Hugo Gilbert, Mohamed Ouaguenouni, Meltem Ozturk, Olivier Spanjaard. Cautious Learning of Multiattribute Preferences. 13th Multidisciplinary Workshop on Advances in Preference Handling, Jul 2022, Vienna, Austria. hal-03694619

HAL Id: hal-03694619

<https://hal.science/hal-03694619v1>

Submitted on 13 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Cautious Learning of Multiattribute Preferences

Hugo Gilbert¹, Mohamed Ouaguenouni², Meltem Ozturk¹, Olivier Spanjaard²

¹Université Paris-Dauphine, Université PSL, CNRS, LAMSADE, 75016 Paris, France

²Sorbonne Université, CNRS, LIP6, 75005, Paris, France

firstname.lastname@lamsade.dauphine.fr,
firstname.lastname@lip6.fr

Abstract

This paper is dedicated to a cautious learning methodology for predicting preferences between alternatives characterized by binary attributes (formally, each alternative is seen as a subset of attributes). By “cautious”, we mean that the model learned to represent the multi-attribute preferences is general enough to be compatible with any strict weak order on the alternatives, and that we allow ourselves not to predict some preferences if the data collected are not compatible with a reliable prediction. A predicted preference will be considered reliable if all the simplest models (following Occam’s razor principle) explaining the training data agree on it. Predictions are based on an ordinal dominance relation between alternatives [Fishburn and LaValle, 1996]. The dominance relation relies on an uncertainty set encompassing the possible values of the parameters of the multi-attribute utility function. Numerical tests are provided to evaluate the richness and the reliability of the predictions made.

1 Introduction

Preference elicitation or preference learning is an important step in setting up a recommender system for a Decision-Maker (DM). It usually consists in querying the DM, e.g. by asking her to assign alternatives to ordered categories. By calling the learning procedure “cautious”, we mean a procedure that complies with two principles that we now describe.

First, the sophistication of the learned multiattribute decision model should be adapted to fit the level of complexity of the stated preferences, hence the choice of a multiattribute utility function f general enough to represent any order \succ of preference, i.e., for any strict weak ordering \succ on a set \mathcal{A} of alternatives, there exists f such that, for any pair $\{A, B\} \subseteq \mathcal{A}$, $f(A) > f(B)$ iff $A \succ B$. In particular, the multi-attribute model we use is able to model positive or negative interactions between attributes [Grabisch *et al.*, 2008].

Second, the predicted pairwise preferences should not depend on the partly arbitrary choice of precise numerical values for the parameters of the model but solely on the stated preferences, hence the design of an *ordinal* learning procedure that maintains an isomorphism between the collected

preferential data and the learned model (in the same spirit as ordinal measurement for problem solving [Bartee, 1971]) by using a polyhedron of possible values for the parameters, reflecting the uncertainty about them. As a consequence of this latter principle, when predicting an unknown pairwise preference between two alternatives A and B , apart from the predictions “ A is preferred to B ” and “ B is preferred to A ”, it is possible that the model does not make a prediction due to a lack of sufficiently rich preferential data (the absence of prediction is preferred to a bad prediction, although a compromise must obviously be made between the reliability of the prediction and the predictive power of the learned model).

Problem. We consider a multiattribute preference elicitation problem, where the attributes are assumed to be binary. Most elicitation procedures make an assumption of a numerical model, defined *a priori*, underlying the DM’s preferences. The originality of our approach consists in allowing the model to be revised during the elicitation process, by modifying the parameters space. The set of model parameters is thus defined *a posteriori* from the preference statements.

A sparse model. Following Fishburn and Lavallo [1996], we consider an underlying numerical model f where the value $f(S)$ of a set S of attributes is an additive combination of parameters, one per subset A of S : $f(S) = \sum_{A \subseteq S: A \in \theta} u_A$. While this model is general enough to model any strict weak ordering on the subsets of attributes, it is inherently intractable as there is a combinatorial set of parameters u_A . To keep a tractable set, similarly to the k -additive variant of this model (which only considers parameters u_A for $|A| \leq k$; Fishburn and Lavallo took $k = 2$), we only consider a restricted family θ of subsets A . We explore different strategies to design θ through the elicitation process. Our goal is to keep it minimal (in a formally defined sense), and yet general enough to fit the training set of pairwise preferences.

Cautious learning. For each pair of alternatives, according to the collected preferential information, our learned model makes a *cautious* prediction: it could either claim which alternative is preferred, or state that the collected information is not sufficient to conclude. In a nutshell, we only make predictions that are consistent with all the simplest models (following Occam’s razor principle) able to explain the stated preferences. The aim is to maximize the ratio of the number of correct preference predictions over the total number of

predictions, while maintaining enough inference power.

Organization of the paper. After giving a brief review of the related work in Section 2, we present the θ -additive utility model in Section 3, as well as the ordinal dominance relation that is inferred if the parameters are only partially specified. In Section 4, we show how to compute the simplest model compatible the collected preferences. Finally, Section 5 is devoted to numerical tests on synthetic preference data.

2 Related work

Preference elicitation, which is part of the broader framework of *preference learning* (see e.g. Fürnkranz and Hüllermeier [2003]), has been studied for a long time in AI, as a preliminary step in any automation of a recommendation task.

We focus here on the elicitation of the parameters of a multiattribute utility function taking into account interactions between attributes (more precisely, learning a partial specification of these parameters yielding a dominance relation between alternatives). In contrast with the setting of active learning which has been widely studied for preference elicitation (see e.g. Guo and Sanner [2010]), we do not assume interactions with a DM but only the knowledge of a “static” training set of examples of pairwise preferences. In this passive learning setting, many classification-based approaches has been proposed, going from perceptrons [Dragone *et al.*, 2017] to Gaussian processes [Chu and Ghahramani, 2005] or Support Vector Machines (SVM) [Domshlak and Joachims, 2005]. These approaches have in common that they consider, as a training set, a set of triples (A, B, c) , where A and B are two alternatives and $c=1$ if $A \succ B$, and $c=0$ otherwise.

A well-known multiattribute utility model that takes into account interactions between attributes, and closely related to the decision model we study in this paper, is the Choquet integral. One of the most recent work about the elicitation of the parameters of a Choquet-related aggregation function integral is that of Bresson *et al.* [2020], in which in particular a perceptron approach is integrated into the learning process of a 2-additive hierarchical Choquet integral [Bresson *et al.*, 2020]. For a broad literature review about learning the parameters of a Choquet integral, the reader may refer to the article by Grabisch *et al.* [2008]. Let us mention in particular the work by Marichal and Roubens [2000], that use a polyhedron to characterize the set of parameters that are compatible with a training set of examples. The idea of defining a polyhedron of uncertainty on the parameters of a utility function goes back at least to the work of Charnetski and Soland [1978]. Their model state that $A \succ B$ if the proportion of parameters that give a better value for A than for B among those that are compatible with the stated preferences is greater than the proportion of parameters that give a better value for B than for A . This principle was also adapted to the case of a Choquet integral by Angilella *et al.* [2015]. In the sequel, we will use a similar polyhedron.

The two works probably closest to our proposal are those of Domshlak and Joachims [2005] and Bigot *et al.* [2012]. For binary attributes, Domshlak and Joachims consider a multiattribute utility function that is a sum of 4^n subutilities over subsets of attribute values and develop an efficient SVM ap-

proach to reveal this utility function, by relying on a kernel method. Bigot *et al.* study the use of generalised additively independent decompositions of utility functions [Fishburn, 1970; Gonzales and Perny, 2005]. They give a polynomial PAC-learner when a constant bound is known on the function’s degree, where the degree is the size of the greatest subset of attributes in the decomposition. Yet, both works do not fit the “cautious learning” framework we consider here.

3 Our Cautious Learning Setting

3.1 Multiattribute Decision Problem

In this paper, we tackle a multiattribute decision problem where alternatives are expressed in the form of a vector of binary attributes. Let $\mathcal{F} = \{a_1, a_2, \dots, a_n\}$ be n binary attributes and $\mathcal{A} \subseteq \{0, 1\}^n$ be the set of alternatives defined on \mathcal{F} . By abuse of notation, for $a_i \in \mathcal{F}$ and $A \in \mathcal{A}$, we will write $a_i \in A$ if the i^{th} component of the vector characterizing A is 1. Moreover, for a subset $S \subseteq \mathcal{F}$ of attributes, we will write $S \subseteq A$ if $a \in A$ for all $a \in S$. For instance, if A corresponds to $(1, 1, 1, 0)$, then $\{a_1, a_3\} \subseteq A$.

We assume that the DM has preferences in the form of a strict weak order over \mathcal{A} . For $A, B \in \mathcal{A}$, we write $A \succ B$ when A is strictly preferred to B , and $A \sim B$ when neither $A \succ B$ nor $B \succ A$ (incomparability).

The aim of preference elicitation is to predict strict pairwise preferences from a training set of examples.

3.2 The θ -additive Model

Cardinal models and additive functions. As the DM’s preferences over \mathcal{A} are modeled as a strict weak order, there exists a real-valued function f such that $\forall A, B \in \mathcal{A}, f(A) > f(B) \Leftrightarrow A \succ B$. Many models assume that f can be represented in a compact way using some sort of additive property.

One of the simplest and most used cardinal models for preference modelling in multiattribute utility theory is the 1-additive model [Keeney *et al.*, 1993]. This model makes the strong assumption that we can find a utility $u(a) \in \mathbb{R}$ for each attribute $a \in \mathcal{F}$ such that for all $A \in \mathcal{A}$, $f(A) = \sum_{a \in A} u(a)$. This assumption is strong because it implies that there is no interaction between the attributes. A weaker assumption is that of *k-additivity* where we suppose the existence of a parameter $u(S) \in \mathbb{R}$ for each $S \in [\mathcal{F}]^k$, where $[\mathcal{F}]^k = \{S \subseteq \mathcal{F} : 1 \leq |S| \leq k\}$. Hence, in the k -additive model, for all $A \in \mathcal{A}$, $f(A) = \sum_{S \in [\mathcal{F}]^k} I_A(S) u_S$, where $I_A(S) = 1$ if $S \subseteq A$ and 0 otherwise, and u_S is an abbreviation for $u(S)$. For example, the 2-additive model makes it possible to account for binary interactions (positive or negative). The n -additive model is general enough to represent *any* strict weak order on \mathcal{A} because it can represent any real-valued set function $f: 2^{\mathcal{F}} \rightarrow \mathbb{R}$ [Grabisch *et al.*, 2000], provided that $f(\emptyset) = 0$. However, it requires to specify $2^n - 1$ parameters. We therefore restrict our attention to additive models requiring fewer parameters.

The θ -additive model. In this paper, we consider a more flexible model which we call the θ -additive model. Given a set $\theta \subseteq 2^{\mathcal{F}}$, and a set function $u: \theta \rightarrow \mathbb{R}$, this model assumes that f is of the form $f(A) = \sum_{S \in \theta} I_A(S) u_S$, where u_S stands again for $u(S)$. In this case, we may also use the notation

$f_{\theta,u}(A)$ instead of $f(A)$. Hence, the 1-additive model is the special case in which θ is $[\mathcal{F}]^1$, and the k -additive model is the special case in which θ is $[\mathcal{F}]^k$.

Example 1. Let $\mathcal{F} = \{a_1, a_2, a_3, a_4\}$ be a set of 4 attributes, $\mathcal{A} = \{0, 1\}^4$ and the preferences of the DM be the strict weak order \succ defined by:

$$\begin{aligned} (0, 1, 1, 1) &\succ (1, 0, 1, 1) \succ (1, 1, 0, 1) \succ (0, 0, 1, 1) \\ &\succ (0, 1, 0, 1) \succ (0, 1, 1, 0) \succ (1, 0, 0, 1) \succ (1, 0, 1, 0) \\ &\succ (1, 1, 0, 0) \succ (0, 0, 0, 1) \succ (0, 0, 1, 0) \succ (0, 1, 0, 0) \\ &\succ (1, 0, 0, 0) \succ (1, 1, 1, 1) \sim (0, 0, 0, 0) \succ (1, 1, 1, 0). \end{aligned}$$

These preferences can be explained by a clear negative interaction when attributes a_1 , a_2 , and a_3 are chosen together (vectors in bold). Interestingly, instead of using a 3-additive model, which would require the definition of 14 parameters, one can use the θ -additive model with $\theta = \{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}, \{a_1, a_2, a_3\}\}$ and $u_{\{a_1\}} = 1$, $u_{\{a_2\}} = 2$, $u_{\{a_3\}} = 3$, $u_{\{a_4\}} = 4$, $u_{\{a_1, a_2, a_3\}} = -10$.

3.3 The Ordinal Dominance Relation

We assume that we only have access to a partial set R of strict pairwise preferences provided by the DM. This set may contain only a few comparisons. Our aim is to use these comparisons (observed preferences) in order to infer other strict pairwise preferences on the set of alternatives. We formalize R as a set of pairs $(A, B) \in \mathcal{A}^2$ such that $A \succ B$.

Moreover, given θ , U_R^θ denotes the set of utility functions on θ that are compatible with the preferences observed in R :

$$U_R^\theta = \{u : \theta \rightarrow \mathbb{R} \mid \forall (A, B) \in R, f_{\theta,u}(A) > f_{\theta,u}(B)\}.$$

Note that, for a given θ , this set U_R^θ can be empty or composed of an infinity of possible utility functions on θ . Notably, if this set is empty then the preferences of the DM cannot be represented by a θ -additive function.

Viewing a θ -additive function as a vector whose dimensions are the subsets S in θ , the set U_R^θ corresponds to the polyhedron defined by the following linear constraints in the $|\theta|$ -dimensional parameter space (where each parameter u_S corresponds to a dimension)¹:

$$\forall (A, B) \in R, \sum_{S \in \theta} I_A(S) u_S - \sum_{S \in \theta} I_B(S) u_S \geq 1. \quad (\text{P1})$$

For a given θ , checking whether or not the preferences of the DM can be represented by a θ -additive function can be evaluated in polynomial time by testing the consistency of the constraints in P1 (e.g., using a linear programming solver).

We denote by Θ_R the set $\{\theta \mid U_R^\theta \neq \emptyset\}$, i.e., the θ 's such that the preferences in R are consistent with a θ -additive function.

Example 2. Coming back to Example 1, setting $\theta = \{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}\}$ yields $U_R^\theta = \emptyset$. In contrast, setting $\theta_1 = \{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}, \{a_1, a_2, a_3\}\}$ yields $U_R^{\theta_1} \neq \emptyset$. In this example, it can be shown² that $\Theta_R = \{\theta : \theta_1 \subseteq \theta\}$.

¹The right hand side of the constraint is here set to 1, but it could be set to any strictly positive constant as utilities u_S are always compatible with R to within a multiplicative factor.

²It has been computer tested by brute force enumeration.

As shown in the previous example, there may be several θ in Θ_R . Moreover, for $\theta \in \Theta_R$, if U_R^θ is compounded of several compatible utility functions, then these utility functions may lead to quite different inferred preferences.

Example 3. Let $\mathcal{F} = \{a_1, a_2, a_3, a_4\}$. Let us assume that, contrary to Example 1, we now only observe preferences on the singletons $\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}$:

$$R = \{((0, 0, 0, 1), (0, 0, 1, 0)), ((0, 0, 1, 0), (0, 1, 0, 0)), ((0, 1, 0, 0), (1, 0, 0, 0))\}$$

The two additive functions u and u' defined by $u(\{a_1\}) = 1$, $u(\{a_2\}) = 2$, $u(\{a_3\}) = 3$, $u(\{a_4\}) = 5$ and $u'(\{a_1\}) = 1$, $u'(\{a_2\}) = 3$, $u'(\{a_3\}) = 4$, $u'(\{a_4\}) = 5$ are both in U_R^θ , but we infer $(1, 0, 0, 1) \succ (0, 1, 1, 0)$ from u while we infer $(0, 1, 1, 0) \succ (1, 0, 0, 1)$ from u' .

This example shows that, given R , choosing a specific function $u \in U_R^\theta$ can lead to infer preferences on the rest of R that are only related to this arbitrary choice and not from the observed preferences [Bartee, 1971]. As we will present in next sections, our aim is to infer preferences for pairs which do not belong to R in a reliable way. In this purpose, we turn to an ordinal model based on the observed preferences which are in R .

Fishburn and Lavalley [1996] showed how one can obtain an ordinal dominance relation from an underlying partially specified 2-additive numerical model. We now explain how the idea can be extended to an underlying θ -additive model.

For a given $\theta \in \Theta_R$, the ordinal dominance relation is denoted by \succ_θ^R , and is independent from the choice of a specific $u \in U_R^\theta$. This binary relation is defined, for each pair A, B in \mathcal{A} , by:

$$A \succ_\theta^R B \Leftrightarrow \forall u \in U_R^\theta, f_{\theta,u}(A) > f_{\theta,u}(B).$$

Naturally, $(A, B) \in R \Rightarrow A \succ_\theta^R B$. Nevertheless, note that binary relation \succ_θ^R is obviously partial, and we define the incomparability relation \sim_θ^R as:

$$A \sim_\theta^R B \Leftrightarrow \exists u, u' \in U_R^\theta, (f_{\theta,u}(A) \geq f_{\theta,u}(B) \text{ and } f_{\theta,u'}(B) \geq f_{\theta,u'}(A)).$$

For any pair A, B of subsets, one can test if $A \succ_\theta^R B$ in polynomial time, by considering the linear program where the objective function $\sum_{S \in \theta} u_S I_B(S) - \sum_{S \in \theta} u_S I_A(S)$ is maximized under constraints P1 (that characterize the set U_R^θ of compatible utility functions). The dominance $A \succ_\theta^R B$ holds iff the optimal value is strictly negative.

If $A \succ_\theta^R B$ then one can predict, based on R and for a θ -additive model, that A is strictly preferred to B . If $A \sim_\theta^R B$ then no prediction is made.

3.4 Sensitivity of the Ordinal Dominance Relation to Changes in R or Θ

We now explore how the relation \succ_θ^R is modified when some new pairwise comparisons are added to R , or removed. Interestingly, adding new pairwise comparisons to R can only enrich binary relation \succ_θ^R , provided the preferences remain representable by a θ -additive function. Conversely, preference $A \succ_\theta^R B$ cannot be reversed by removing pairwise comparisons from R . More formally:

Proposition 1. *Given a set R of strict pairwise comparisons, and $\theta \in \Theta_R$, if $R' \subseteq R$, then we have: (i) $\theta \in \Theta_{R'}$; (ii) $A \succ_{\theta}^{R'} B \Rightarrow A \succ_{\theta}^R B$; (iii) $A \succ_{\theta}^R B \Rightarrow \neg(B \succ_{\theta}^{R'} A)$.*

We now study how the relation \succ_{θ}^R is modified when θ is restricted or extended. If θ is restricted, then the relation \succ_{θ}^R can only be enriched. Conversely, if θ is extended, then a preference $A \succ_{\theta}^R B$ cannot be reversed after the extension.

Proposition 2. *For $\theta, \theta' \in \Theta_R$, if $\theta' \subseteq \theta$, then we have: (i) $A \succ_{\theta}^R B \Rightarrow A \succ_{\theta'}^R B$; (ii) $A \sim_{\theta'}^R B \Rightarrow A \sim_{\theta}^R B$; (iii) $A \succ_{\theta'}^R B \Rightarrow \neg(B \succ_{\theta}^R A)$.*

Note that many different θ -additive models may be compatible with the collected preferences in R . In particular, if a θ' -additive model is compatible with R , then any θ -additive model such that θ extends θ' is also compatible with R . A natural way to decide which θ -additive models to consider is to follow the inclusion relationship on Θ_R , by considering the sets θ that are minimal w.r.t. inclusion. For computational efficiency, we will use a refinement of the inclusion relationship, that we detail in the next section.

4 The Minimal Compatible Models and The Unifying Model

Note that there always exists a θ able to represent R ; at worst, we can put all the subsets of \mathcal{F} in θ . Our choice of a specific θ among the various ones that yield a θ -additive model able to explain the collected preferences in R is guided by two criteria, namely:

- First, following the philosophical principle of parsimony that the simpler of two explanations is to be preferred (Occam's razor [Blumer *et al.*, 1987]), we consider subsets θ that minimize the complexity of interactions between the attributes; to measure this complexity, we use the *degree* of $f_{\theta,u}$, namely $\max\{|S| : S \in \theta\}$ (i.e., the greatest cardinality of a subset of interacting attributes).
- Second, if two different θ have the same *degree*, we prefer the one having *sparsest* representation [Zhang *et al.*, 2015], i.e., the one which minimizes $|\theta|$ (which corresponds to the number of non-zero parameters u_S).

This two criteria define a lexicographic binary relation on Θ_R , refining \subseteq and denoted by \sqsubseteq_{lex} . We call $\theta \in \Theta_R$ which are minimal according to \sqsubseteq_{lex} , *simplest* θ of R and we denote by Θ_R^{\min} their set: $\Theta_R^{\min} = \{\theta \in \Theta_R \mid \nexists \theta', \theta' \sqsubseteq_{lex} \theta\}$.

Note that sometimes the simplest model may contain more elements than another model which has a bigger degree:

Example 4. *Let $R = \{(1, 1, 0, 0) \succ (0, 0, 1, 1), (1, 1, 0, 0) \succ (1, 0, 1, 0)\}$. It is easy to see that we can find a θ with one element containing a subset of cardinality 2 ($\theta = \{\{a_1, a_2\}\}$), however we will prefer having a θ consistent with a 1-additive model even if there are more elements in it: $\theta' = \{\{a_1\}, \{a_2\}\}$ or $\theta'' = \{\{a_1\}, \{a_3\}\}$ or $\theta''' = \{\{a_2\}, \{a_3\}\}$.*

4.1 Computation of Θ_R^{\min}

To compute the set Θ_R^{\min} from R , we perform an enumeration of all possible minimal θ sets by using Algorithm 1 called with $\Theta_R^{\min} = \theta = \emptyset$, and $\underline{\theta} = 2^{\mathcal{F}}$. The parameters used by

Algorithm 1 are the list Θ_R^{\min} under construction, a representative $\underline{\theta}$ of Θ_R^{\min} used to test whether θ is minimal w.r.t. \sqsubseteq_{lex} , the current θ under examination (i.e., whose membership to Θ_R^{\min} is being guessed) and the set R of collected preferences.

To perform this enumeration, we rely on :

- a depth first search strategy, where each node corresponds to a possible θ , the root is initialized with $\theta = \emptyset$, and a node is expanded by investigating the possible sets S that may break (i.e., invalidate) the certificate I that R is not compatible with a θ -additive model (lines 8 to 11 in Algorithm 1); we explain below how a certificate I is defined and determined.
- a pruning strategy consisting in exploring only nodes who correspond to sets θ that are not dominated by the ones in Θ_R^{\min} w.r.t. \sqsubseteq_{lex} (lines 2-3, 10 in Algorithm 1).

Algorithm 1 BuildThetaMin($\Theta_R^{\min}, \underline{\theta}, \theta, R$)

```

1: if  $R$  can be represented by a  $\theta$ -additive model then
2:   if  $\theta \sqsubseteq_{lex} \underline{\theta}$  then
3:      $\Theta_R^{\min} \leftarrow \{\theta\}$ ;
4:      $\underline{\theta} \leftarrow \theta$ ;
5:   else
6:      $\Theta_R^{\min} \leftarrow \Theta_R^{\min} \cup \{\theta\}$ ;
7:   else
8:     Find certificate  $I$  and preference set  $C$  by solving  $\mathcal{D}_{\theta}$ ;
9:     for  $S \in \{T \subset A \setminus B : (A, B) \in C \text{ or } (B, A) \in C\}$  do
10:      if  $S$  breaks certificate  $I$  and not  $\theta \sqsubseteq_{lex} \theta \cup \{S\}$  then
11:        BuildThetaMin( $\Theta_R^{\min}, \underline{\theta}, \theta \cup \{S\}, R$ );

```

Determining if R can be represented by a θ -additive model (line 1 of Algorithm 1). Given a parameter set θ , the following linear program \mathcal{P}_{θ} , where there is one positive variable $e_{A,B}$ for each pair (A, B) in R , and one free variable u_S for each set S in θ , determines if the set R of observed strict preferences can be represented by a θ -additive model:

$$\begin{aligned}
(\mathcal{P}_{\theta}) \quad & \min_{e_{A,B}, u_S} \sum_{(A,B) \in R} e_{A,B} \\
& \sum_{S \in \theta} (I_A(S) - I_B(S)) u_S \geq 1 - e_{A,B}, \quad \forall (A, B) \in R \\
& e_{A,B} \geq 0, \quad \forall (A, B) \in R
\end{aligned}$$

The preferences in R can be represented by a θ -additive model if the optimal value of \mathcal{P}_{θ} is 0. Indeed, in this case we can find values for variables u_S that respect all the preferences in R without the help of the additional slack variables $e_{A,B}$.

Program \mathcal{P}_{θ} is probably the most intuitive program to test if R can be represented by the θ -additive model. However, we will work instead on its dual \mathcal{D}_{θ} :

$$\begin{aligned}
(\mathcal{D}_{\theta}) \quad & \max_{\lambda_{A,B}} \sum_{(A,B) \in R} \lambda_{A,B} \\
& \sum_{(A,B) \in R} (I_A(S) - I_B(S)) \lambda_{A,B} = 0, \quad \forall S \in \theta \\
& 0 \leq \lambda_{A,B} \leq 1, \quad \forall (A, B) \in R
\end{aligned}$$

If the optimal value of \mathcal{D}_θ is strictly positive, we must add at least another set to θ to represent the preferences in R .

Finding a certificate (line 9 of Algorithm 1). Let $I = (\lambda_{A,B}^* : (A, B) \in R)$ be an optimal solution to program (\mathcal{D}_θ) such that $\sum_{(A,B) \in R} \lambda_{A,B}^* > 0$. Note that the values in I make it possible to identify a set of preferences $C = \{(A, B) : \lambda_{A,B}^* > 0\}$ that cannot be represented by the current θ -additive model, and that I is in some sense a certificate for the incapacity to represent C and thus R (because $C \subseteq R$). In this case, one should add a set T to θ . This amounts to adding the constraint $\sum_{(A,B) \in R} (I_A(T) - I_B(T))\lambda_{A,B} = 0$ to \mathcal{D}_θ . Importantly, note that this may only decrease the optimal value of (\mathcal{D}_θ) if $\sum_{(A,B) \in R} (I_A(T) - I_B(T))\lambda_{A,B}^* \neq 0$. Hence, the different candidates to add to θ will be precisely the sets T that satisfy this condition. When adding such a set to θ we will informally say that we break I ³.

Finding a set S breaking I (lines 10-14 of Algorithm 1). Note that a set S breaking I can always be found (even efficiently) as R can be represented by any θ -additive model with $\{A, B : (A, B) \in R\} \subseteq \theta$. Hence, a set S breaking I can always be found in $\{A, B : (A, B) \in R\}$. However, to keep θ “simple” we explore more systematically the sets that can break I in order to find simple ones. In a nutshell, we enumerate all the sets in $\{S \subset A \setminus B : (A, B) \in C \text{ or } (B, A) \in C\}$. Indeed, each of these subsets may change the scores of sets appearing in C and hence break the certificate I .

4.2 The Unifying Model

Instead of predicting $A \succ B$ if $A \succ_\theta^R B$ for all $\theta \in \Theta_R^{\min}$, we consider a single set θ “synthesizing” Θ_R^{\min} and infer preferences from it, because they are more easily explainable. An intuitive idea consists of taking the union of all the simplest θ . We call this model *unifying model* and denote it by θ_R^* :

$$\theta_R^* = \cup_{\theta \in \Theta_R^{\min}} \theta.$$

Using the unifying model, we guarantee not to contradict the preferences that are compatible with all the θ in Θ_R^{\min} .

Proposition 3. *Let R be the set of observed preferences on the elements of \mathcal{A} , let Θ_R^{\min} be the set of simplest θ -models compatible with R and $\theta_R^* = \cup_{\theta \in \Theta_R^{\min}} \theta$, then $\forall \theta \in \Theta_R^{\min}, \forall A, B \in \mathcal{A}$*

$$A \succ_{\theta_R^*}^R B \Rightarrow A \succ_\theta^R B.$$

Unfortunately the inverse is not true, i.e., it is possible that $A \succ_\theta^R B$ for $\theta \in \Theta_R^{\min}$ but not $A \succ_{\theta_R^*}^R B$. Example 7 in appendix illustrates this point.

5 Numerical Tests

Numerical tests were carried out on Google Colab (2 virtual CPU at 2.2GHz, 13GB RAM). The objective of these tests is twofold: 1) evaluating the *accuracy rate* of the predictions, namely the number of correct pairwise preference predictions over the total number of predicted preferences, if the set θ is *known* beforehand; 2) evaluating the same metric if the set θ is *unknown* beforehand and learned with Algorithm 1.

³This can be thought of as solving a separation problem, by providing an hyperplane separating I from the polytope of $\mathcal{D}_{2\mathcal{F}}$.

5.1 The Tier List Framework

We place ourselves in an elicitation context where each query consists in asking the DM to position an alternative in a tier list of ordered classes (i.e., the worst alternatives in category 1, the second worst alternatives in category 2, etc.). Formally, we assume that the user gives us access to a function $\gamma : \mathcal{A} \rightarrow \mathbb{N}$ that associates each alternative to a class in the tier list such that $\gamma(A) > \gamma(B) \Rightarrow A \succ B$. Note that $\gamma(A) = \gamma(B)$ does not mean here that A and B are indifferent, but that the user does not know how to compare them.

Positioning one alternative in the tier list allows us to inter-actively collect numerous strict pairwise preference relations while keeping a low cognitive burden compared to asking for pairwise comparisons or for scores (one score per alternative).

5.2 Synthetic Generation of a Tier List

This section details our simulation of the creation of a tier list from a θ -additive function modeling the DM’s preferences.

Sampling a θ -additive Function $f_{\theta,u}$

For sampling a function $f_{\theta,u}$, we first sample a set θ and then sample parameters u_S for $S \in \theta$. More precisely, the generation of θ is achieved as follows. First, θ is initialised as the set of singletons $\{a_1\}, \{a_2\}, \dots, \{a_n\}$, then we add $\lfloor \alpha \times (2^{|\mathcal{F}|} - |\mathcal{F}|) \rfloor$ subsets of attributes, where the coefficient $\alpha \in [0, 1]$ makes it possible to control the model’s complexity: for $\alpha = 0$, only the singletons are in θ , which yields the simple additive utility model, and for $\alpha = 1$, all subsets of attributes are present, with yields the most general utility model. Each subset S is sampled according to a parameter $p \in (0, 1]$:

1. Initialize S as a singleton by uniformly sampling in \mathcal{F} .
2. Uniformly sample another attribute in \mathcal{F} and add it to S .
3. Exit this process if $S = \mathcal{F}$.
4. Exit this process with a probability p otherwise go to 2.

The expected size of each S we add can be approximated by:

$$\mathbb{E}[|S|] = 2 + (1 - p - (1 - p)^{n-1})/p.$$

Table 1 gives some hint of the expected size of each S according to p . Once θ is set, we sample the parameters u_S for each $S \in \theta$ with a normal distribution $\mathcal{N}(0, \sigma)$. The sampling of $f_{\theta,u}$ thus depends on three parameters p , α and σ . In the tests, p varies in $[0.1, 0.9]$, α in $[0.1, 0.5]$, and we set $\sigma = 100$.

p	0.2	0.4	0.6	0.8	1
$\mathbb{E}[S]$	3.95	3.18	2.62	2.25	2.00

Table 1: Expected size of subsets S w.r.t. p .

Example 5. *If $n = 4$, $p = 0.3$, $\alpha = 0.1$, then $\lfloor 0.1(2^5 - 5) \rfloor = 2$ subsets S are sampled in addition to the singletons. This may yield the parameter values given in Table 2.*

Subset	Value	Subset	Value
{0}	148.85	{4}	191.00
{1}	186.75	{1,3,4}	-26.80
{2}	90.60	{0,2}	80.24
{3}	-86.12		

Table 2: Example of parameter values.

From $f_{\theta,u}$ to a Tier List

The function $\gamma : \mathcal{A} \rightarrow \mathbb{N}$ that simulates the user assignment of alternatives into a tier list, called tier function hereafter, relies on a parameter t representing the number of categories. The range of scores $f_{\theta,u}(A) = \sum_{S \in \theta} u_S I_A(S)$ of alternatives A is partitioned into t equally-sized intervals between the min score $f_0 = \min_{A \in \mathcal{A}} f_{\theta,u}(A)$ and the max score $f_t = \max_{A \in \mathcal{A}} f_{\theta,u}(A)$. The function γ is then defined by:

$$\gamma(A) = \min\{1 \leq k \leq t : f_{\theta,u}(A) \leq u_k\}.$$

Put another way, we associate to each subset the interval where its utility lies. In general, the more categories we add, the less incomparabilities we will have (alternatives assigned to the same category), but the user will have to make more efforts to assign the alternatives to categories.

Example 6. *Coming back to Example 5, let $\mathcal{A} = \{0, 1\}^n$. Then $\max_{A \in \mathcal{A}} f_{\theta,u}(A) = 616.41$ and $\min_{A \in \mathcal{A}} f_{\theta,u}(A) = -86.12$. Assume that one partitions into $t = 3$ categories. The intervals are then $[-86.12, 148.05]$, $(148.05, 382.23]$ and $(382.23, 616.41]$. Subset $\{1, 2, 3\}$ is then assigned to category 2 because its utility 191.23 belongs to $(148.05, 382.23]$.*

5.3 Baseline Models

In the following, the ordinal model studied in the paper is denoted by ORD. In this part, we will briefly introduce the baseline models to which ORD is compared.

Linear Programming Model (LPM). As a first baseline model, we compare our approach with the model consisting in setting parameters u_S at their optimal values for the linear program \mathcal{P}_θ of page 4, and predicting that $A \succ B$ if $f_{\theta,u}(A) > f_{\theta,u}(B)$. In the experiments, if θ is known beforehand, only the constraints set of \mathcal{P}_θ grows, while if θ is unknown, both the variables and the constraints may change when R grows.

Support Vector Machine (SVM). This baseline model is inspired by an approach proposed by Domshlak and Joachims [2005]. An SVM approach is a supervised learning method for binary classification: each example in the dataset is labeled by 0 or 1; an SVM is learned from the dataset, from which labels are inferred for new examples. In our setting, each preference $A \succ B$ in R yields two examples: a $(2m+1)$ -dimensional vector $(v_A^\theta, v_B^\theta, 1)$ and another vector $(v_B^\theta, v_A^\theta, 0)$. That is, the third component of $(v_A^\theta, v_B^\theta, c)$ is $c = 1$ if A is preferred to B , and $c = 0$ if it is not. Note that, when inferring labels (and thus predicting preferences), it may happen that (v_A^θ, v_B^θ) and (v_B^θ, v_A^θ) get the same label (0 or 1). In this case, no strict preference is predicted.

5.4 Experiment with a Known θ

In the first experiment, we compared the two above baseline models with our ordinal model when the θ used to generate the tier function γ is known beforehand.

Used metrics. To evaluate the accuracy of each model, we rely on the following measures:

- Correct answers (C): an inferred preference $A \succ B$ is said to be correct if $\gamma(A) > \gamma(B)$.
- Wrong answers (W): an inferred preference $A \succ B$ is said to be wrong if $\gamma(A) < \gamma(B)$.

Given a model (ORD, LPM or SVM), a preference between A and B is *inferred* if the preference is *not* already present in R and the model states that $A \succ B$ or $B \succ A$ (but not both). We denote by T the total number of inferred preferences. Note that $C + W \leq T$ because it may happen that $\gamma(A) = \gamma(B)$. The Absolute Correct Rate (ACR) is defined from T and C :

$$ACR = C/T$$

Experimental setting. The experiment was conducted with $|\mathcal{F}| = 5$, $t = 12$, $\sigma = 100$, and two sets of parameters (α, p) , namely $(\alpha, p) = (0.1, 0.9)$ and $(\alpha, p) = (0.3, 0.7)$. Roughly speaking, the former set of parameters generates tier functions with low interactions, while the latter generates tier functions with high interactions. For each couple (α, p) , we sample three random tier functions and, for each one, we train each model with a budget of 25 assignments to categories. The test examples are generated as follows: we randomly sample 10 alternatives A_1, \dots, A_{10} in \mathcal{A} and we consider all pairs $\{A_i, A_j\}$ for $i \neq j$. We count the number T of inferred preferences for these pairs, and we evaluate the ACR. To smooth the results, they are averaged over 10 different tier functions, and 5 samples of ten alternatives for each of them.

Results and discussion. The results are presented in Figures 1 and 2, where the x-axis gives the size of the training set and the curves show the mean and 95% confidence interval. The curves show how the average number of inferred preferences and the average ACR evolve with the size of the training set (from 1 to 25 assignments of alternatives to categories). In both figures, we see that the number of inferred preferences grows more slowly with ORD than with LPM and SVM, in accordance with the principle of cautious learning. However, the accuracy is better, as reflected by the curve of ACR for ORD that is consistently above the curves obtained for LPM and SVM. As one could expect, when the interactions are high (Figure 2), and thus the number of parameters u_S is significant, a larger learning set is required to make it possible to infer numerous pairwise preferences with ORD. Note that, when the number of assignments available in the training set is low, the confidence interval for the curve of ACR for ORD is wide. This is related to the fact that few preferences are inferred and therefore a wrong prediction drastically change the ACR. However, after 15 assignments, the number of inferred preferences becomes higher, and the ACR for ORD outperforms the ACR for LPM and SVM. Comparing Figure 1 and Figure 2, we can even see that, after 25 assignments, the difference in ACR is greater with high interactions than with low interactions. We ascribe this to the fact that the three models behave similarly with low numbers of parameters u_S ($|\theta|$ not far from n) because the polyhedron of compatible utilities is small. We also notice in the two figures that the number of inferred preferences is always greater with LPM and SVM than with ORD. Put another way, ORD represents a different trade-off between the number of preferences that can be predicted and their accuracy.

5.5 Experiment with an Unknown θ

In this section, we investigate the behavior of the models when θ is learned at the same time as parameters u_S ($\forall S \in \theta$).

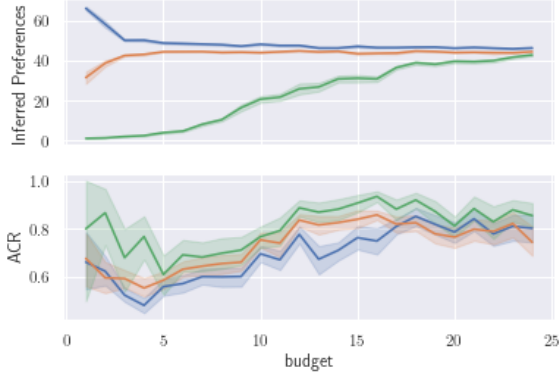


Figure 1: T (top) and ACR (bottom) with ORD (green), LPM (blue) and SVM (orange) for a known θ and $(\alpha, p, t) = (0.1, 0.9, 12)$.



Figure 3: T (top) and ACR (bottom) with ORD (green), LPM (blue) and SVM (orange) for an unknown θ and $(\alpha, p, t) = (0.1, 0.9, 9)$.



Figure 2: T (top) and ACR (bottom) with ORD (green), LPM (blue) and SVM (orange) for a known θ and $(\alpha, p, t) = (0.3, 0.7, 12)$.



Figure 4: T (top) and ACR (bottom) with ORD (green), LPM (blue) and SVM (orange) for an unknown θ and $(\alpha, p, t) = (0.3, 0.7, 9)$.

Experimental setting. The experimental setting is similar to the previous one, except that the number of categories in the tier lists is set to $t = 9$. For all models (ORD, LPM and SVM), the set θ is updated after each assignment of an alternative to a category, by using Algorithm 1.

Results and discussion. The results are presented in Figures 3 and 4, with the same conventions as above. Similarly to the case of a known θ , we see that model ORD outperforms models LPM and SVM in terms of accuracy. We notice small irregularities in the inferred preferences curve of ORD, due to the fact that ORD infers less preferences each time θ is updated because the polyhedron of compatible parameters expands when dimensions are added (corresponding to new subsets in θ). Figure 5 shows the result of another experiment where the models are trained twice: once using the actual θ used to generate the synthetic preferences in R , and a second time using the θ obtained by computing a unifying model (see Section 4.2). Interestingly, both learning curves are close to each other, which tends to show that the learned θ is relevant.

6 Conclusion

We have presented here a “cautious” method for learning pairwise multiattribute preferences. The model we use is not restrictive, in the sense that any preference relation on the space of alternatives can be represented. The learning method



Figure 5: Number of inferred preferences and ACR with the real θ (in orange) and with θ_R^* (in blue), for $(\alpha, p, t) = (0.3, 0.7, 9)$.

achieves a trade-off between the number of predicted preferences and the accuracy of the predictions, by relying on an ordinal dominance relation between alternatives.

Several research directions are worth investigating, among which the adaptation of the approach to an active learning setting where one interactively determines a sequence of queries to minimize the cognitive burden for a DM, or the examination of other definitions of the set Θ_R^{min} of simplest models compatible with R .

Acknowledgements

We acknowledge a financial support from the project THEMIS ANR20-CE23-0018 of the French National Research Agency (ANR).

References

- [Angilella *et al.*, 2015] Silvia Angilella, Salvatore Corrente, and Salvatore Greco. Stochastic multiobjective acceptability analysis for the choquet integral preference model and the scale construction problem. *European Journal of Operational Research*, 240(1):172–182, January 2015.
- [Bartee, 1971] Edwin M Bartee. Problem solving with ordinal measurement. *Management Science*, 17(10):B-622, 1971.
- [Bigot *et al.*, 2012] Damien Bigot, Hélène Fargier, Jérôme Mengin, and Bruno Zanuttini. Using and learning gaidecompositions for representing ordinal rankings. In *ECAI'2012 workshop on Preference Learning (PL 2012)*, pages 5–10. Johannes Fürnkranz and Eyke Hüllermeier, 2012.
- [Blumer *et al.*, 1987] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Occam’s razor. *Information processing letters*, 24(6):377–380, 1987.
- [Bresson *et al.*, 2020] Roman Bresson, Johanne Cohen, Eyke Hüllermeier, Christophe Labreuche, and Michèle Sebag. Learning 2-additive hierarchical choquet integrals with non-monotonic utilities. In *DA2PL 2020*, 2020.
- [Charnetski and Soland, 1978] Johnnie R Charnetski and Richard M Soland. Multiple-attribute decision making with partial information: the comparative hypervolume criterion. *Naval Research Logistics Quarterly*, 25(2):279–288, 1978.
- [Chu and Ghahramani, 2005] Wei Chu and Zoubin Ghahramani. Preference learning with gaussian processes. In *ICML-05*, pages 137–144, 2005.
- [Domshlak and Joachims, 2005] Carmel Domshlak and Thorsten Joachims. Unstructuring user preferences: efficient non-parametric utility revelation. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, pages 169–177, 2005.
- [Dragone *et al.*, 2017] Paolo Dragone, Stefano Teso, and Andrea Passerini. Constructive preference elicitation over hybrid combinatorial spaces. *CoRR*, abs/1711.07875, 2017.
- [Fishburn and Lavalley, 1996] P. C. Fishburn and I. H. Lavalley. Binary interactions and subset choice. *European Journal of Operational Research*, 92:182–192, 1996.
- [Fishburn, 1970] Peter C Fishburn. Utility theory for decision making. Technical report, Research analysis corp McLean VA, 1970.
- [Fürnkranz and Hüllermeier, 2003] Johannes Fürnkranz and Eyke Hüllermeier. Pairwise preference learning and ranking. In *European conference on machine learning*, pages 145–156. Springer, 2003.
- [Gonzales and Perny, 2005] Christophe Gonzales and Patrice Perny. GAI Networks for Decision Making under Certainty. In *19th International Joint Conference on Artificial Intelligence – workshop on advances in preference handling*, pages 100–105, Edinburgh, United Kingdom, 2005.
- [Grabisch *et al.*, 2000] Michel Grabisch, Jean-Luc Marichal, and Marc Roubens. Equivalent representations of set functions. *Mathematics of Operations Research*, 25(2):157–178, 2000.
- [Grabisch *et al.*, 2008] Michel Grabisch, Ivan Kojadinovic, and Patrick Meyer. A review of methods for capacity identification in choquet integral based multi-attribute utility theory: Applications of the kappalab r package. *European journal of operational research*, 186(2):766–785, 2008.
- [Guo and Sanner, 2010] Shengbo Guo and Scott Sanner. Multiattribute bayesian preference elicitation with pairwise comparison queries. In *International Symposium on Neural Networks*, pages 396–403. Springer, 2010.
- [Keeney *et al.*, 1993] Ralph L Keeney, Howard Raiffa, and Richard F Meyer. *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press, 1993.
- [Marichal and Roubens, 2000] Jean-Luc Marichal and Marc Roubens. Determination of weights of interacting criteria from a reference set. *European journal of operational Research*, 124(3):641–650, 2000.
- [Zhang *et al.*, 2015] Zheng Zhang, Yong Xu, Jian Yang, Xuelong Li, and David Zhang. A survey of sparse representation: algorithms and applications. *IEEE access*, 3:490–530, 2015.

Appendix

Proposition 1. Given a set R of strict pairwise comparisons, and $\theta \in \Theta_R$, if $R' \subseteq R$, then we have: (i) $\theta \in \Theta_{R'}$; (ii) $A \succ_{\theta}^{R'} B \Rightarrow A \succ_{\theta}^R B$; (iii) $A \succ_{\theta}^R B \Rightarrow \neg(B \succ_{\theta}^{R'} A)$.

Proof. (i) If all the preferences in R can be represented by a θ -additive function, then so can the preferences in R' as R' is compounded of a subset of the preferences in R .

(ii) If the preferences in R' imply that A should be necessarily strictly preferred to B , then R will imply the same condition as R also contains the same preference constraints as in R' .

(iii) The contrapositive is proved as follows: $B \succ_{\theta}^{R'} A \Rightarrow B \succ_{\theta}^R A$ by (ii), and $B \succ_{\theta}^R A \Rightarrow \neg(A \succ_{\theta}^{R'} B)$ because strict preferences are asymmetrical. \square

Proposition 2. For $\theta, \theta' \in \Theta_R$, if $\theta' \subseteq \theta$, then we have: (i) $A \succ_{\theta}^R B \Rightarrow A \succ_{\theta'}^R B$; (ii) $A \sim_{\theta}^R B \Rightarrow A \sim_{\theta'}^R B$; (iii) $A \succ_{\theta'}^R B \Rightarrow \neg(B \succ_{\theta}^R A)$.

Proof. (i) is true because if $f_{\theta,u}(A) > f_{\theta,u}(B)$ for all $u \in U_R^{\theta}$, then we should also have $f_{\theta',u}(A) > f_{\theta',u}(B)$ for all $u \in U_R^{\theta'}$. Indeed, each element of $U_R^{\theta'}$ can be seen as a utility function in U_R^{θ} in which the parameters u_S are set to 0 for $S \in \theta \setminus \theta'$.

(ii) follows by a similar argument as for (i).

(iii) The contrapositive is proved as follows: $B \succ_{\theta}^R A \Rightarrow B \succ_{\theta'}^R A$ by (i), and $B \succ_{\theta'}^R A \Rightarrow \neg(A \succ_{\theta}^R B)$ because strict preferences are asymmetrical. \square

Example 7. Let's take $\mathcal{A} = \{0, 1\}^4$ ($\mathcal{F} = \{a_1, a_2, a_3, a_4\}$) and observed preferences R as in the following:

$$(1, 1, 1, 0) \succ (0, 0, 0, 1) \succ \emptyset \succ (0, 1, 1, 0).$$

We have $\Theta_R^{\min} = \{\theta_1, \theta_2\}$ with $\theta_1 = \{a_1, a_3, a_4\}$, $\theta_2 = \{a_1, a_2, a_4\}$, and thus $\theta^* = \{a_1, a_2, a_3, a_4\}$.

The polyhedron resulting from θ_1 is:

$$\begin{aligned} u_1 + u_3 &> u_4 \\ u_1 + u_3 &> 0 \\ u_1 &> 0 \\ u_4 &> 0 \\ u_4 &> u_3 \\ 0 &> u_3 \end{aligned}$$

From $u_3 < 0$ and $u_1 + u_3 > 0$ it results that $u_1 > 0$ and since $u_2 = 0$ because $a_2 \notin \theta_1$ we have

$$(1, 0, 0, 0) \succ_{\theta_1}^R (0, 1, 0, 0).$$

The polyhedron resulting from θ_2 is:

$$\begin{aligned} u_1 + u_2 &> u_4 \\ u_1 + u_2 &> 0 \\ u_1 &> 0 \\ u_4 &> 0 \\ u_4 &> u_2 \\ 0 &> u_2 \end{aligned}$$

From $u_2 < 0$ and $u_1 > 0$ we have

$$(1, 0, 0, 0) \succ_{\theta_2}^R (0, 1, 0, 0).$$

Hence, $(1, 0, 0, 0)$ is strictly preferred to $(0, 1, 0, 0)$ for both $\succ_{\theta_1}^R$ and $\succ_{\theta_2}^R$.

Yet, the polyhedron resulting from θ^* is

$$\begin{aligned} u_1 + u_2 + u_3 &> u_4 \\ u_1 + u_2 + u_3 &> 0 \\ u_1 &> 0 \\ u_4 &> 0 \\ u_4 &> u_2 + u_3 \\ 0 &> u_2 + u_3 \end{aligned}$$

And we can verify that

$$u = \{u_1 = 3; u_2 = 5; u_3 = -6; u_4 = 1\} \in U_R^{\theta^*}$$

and since $f_{\theta^*,u}(0, 1, 0, 0) > f_{\theta^*,u}(1, 0, 0, 0)$ the preference $(1, 0, 0, 0) \succ_{\theta^*}^R (0, 1, 0, 0)$ does not hold while $(1, 0, 0, 0)$ is strictly preferred to $(0, 1, 0, 0)$ for both $\succ_{\theta_1}^R$ and $\succ_{\theta_2}^R$.