



HAL
open science

Robust ordinal regression for subsets comparisons with interactions

Hugo Gilbert, Mohamed Ouaguenouni, Meltem Öztürk, Olivier Spanjaard

► **To cite this version:**

Hugo Gilbert, Mohamed Ouaguenouni, Meltem Öztürk, Olivier Spanjaard. Robust ordinal regression for subsets comparisons with interactions. *European Journal of Operational Research*, In press, 320 (1), pp.146-159. 10.1016/j.ejor.2024.07.021 . hal-04677617

HAL Id: hal-04677617

<https://hal.science/hal-04677617v1>

Submitted on 18 Sep 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License



HAL
open science

Robust ordinal regression for subsets comparisons with interactions

Hugo Gilbert, Mohamed Ouaguenouni, Meltem Öztürk, Olivier Spanjaard

► **To cite this version:**

Hugo Gilbert, Mohamed Ouaguenouni, Meltem Öztürk, Olivier Spanjaard. Robust ordinal regression for subsets comparisons with interactions. *European Journal of Operational Research*, 2025, 320, pp.146 - 159. 10.1016/j.ejor.2024.07.021 . hal-04700036

HAL Id: hal-04700036

<https://hal.science/hal-04700036>

Submitted on 17 Sep 2024

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.



Decision support

Robust ordinal regression for subsets comparisons with interactions[☆]Hugo Gilbert^b, Mohamed Ouaguenouni^{a,*}, Meltem Öztürk^{a,b,1}, Olivier Spanjaard^a^a Sorbonne Université, CNRS, LIP6, Paris, F-75005, France^b Université Paris Dauphine, PSL University, CNRS, LAMSADE, Paris, F-75016, France

ARTICLE INFO

Keywords:

Robust ordinal regression
 Preference elicitation
 Positive and negative interactions
 Subsets comparisons

ABSTRACT

This paper is devoted to a robust ordinal method for learning the preferences of a decision maker between subsets. The decision model, derived from Fishburn and LaValle (1996) and whose parameters we learn, is general enough to be compatible with any strict weak order on subsets, thanks to the consideration of possible interactions between elements. Moreover, we accept not to predict some preferences if the available preference data are not compatible with a reliable prediction. A predicted preference is considered reliable if all the simplest models (Occam's razor) explaining the preference data agree on it. Following the robust ordinal regression methodology, our predictions are based on an uncertainty set encompassing the possible values of the model parameters. We define a new ordinal dominance relation between subsets and design a procedure to determine whether this dominance relation holds. Numerical tests are provided on synthetic and real-world data to evaluate the richness and reliability of the preference predictions made.

1. Introduction

Preference elicitation (or preference learning) is an important step in setting up a recommender system for decision making. In this preference elicitation setting, our focus is on determining the parameters of a Decision Maker (DM) over subsets, by comparing subsets of elements. The preferences are depicted using a highly adaptable model whose versatility stems from its ability to incorporate positive or negative synergies between elements (Grabisch, Kojadinovic, & Meyer, 2008). Moreover, we provide an ordinally robust approach, in the sense that the preferences we infer do not rely on arbitrarily specified parameter values, but on the set of all parameter values that are compatible with the observed preferences. Importantly, another distinctive feature of our approach is its ability to learn the parameter set itself (not only the values of parameters).

The preference model we consider can be used in different contexts, depending on the nature of the subsets we are comparing. The subsets are represented by binary vectors, showing the presence or absence of an element in the subset. The elements of a subset can be for example:

- individuals (in the comparison of coalitions, teams, etc.),

- binary attributes (in the comparison of multiattribute alternatives),
- objects (in the comparison of subsets in a subset choice problem), etc.

For illustration, a toy example of such an elicitation context could be a coffee shop trying to determine its customers' favorite frozen yogurt flavor combination by offering them to test a small number of flavor combinations rather than having them taste each combination.

Objective of the paper. Our objective is to design a preference elicitation procedure that complies with the two following principles.

First, the sophistication of the learned preference model should be able to fit any level of complexity of the stated preferences. For this purpose, we use a utility function f general enough to represent any order $>$ of preference, i.e., for any strict weak order $>$ on a set \mathcal{A} of alternatives (i.e., subsets) there exists f such that, for any pair $\{A, B\} \subseteq \mathcal{A}$, $f(A) > f(B)$ iff $A > B$. Note that we also aim to make the model as simple as possible, in the sense that the parameter set remains as concise as possible (*sparse model*). Second, the predicted pairwise preferences should not depend on the partly arbitrary choice of precise numerical values for the parameters of the model

[☆] This paper is a revised and extended version of a workshop paper at MPREF 2022, and an extended abstract at AAMAS 2023: H. Gilbert, M. Ouaguenouni, M. Öztürk, O. Spanjaard, *Cautious Learning of Multiattribute Preferences*, 13th Workshop MPREF, Jul 2022, Vienna, Austria. H. Gilbert, M. Ouaguenouni, M. Öztürk, O. Spanjaard, *Robust Ordinal Regression for Collaborative Preference Learning with Opinion Synergies*, AAMAS 2023, pp. 2439-2441.

* Corresponding author.

E-mail addresses: hugo.gilbert@lamsade.dauphine.fr (H. Gilbert), mohamed.ouaguenouni@lip6.fr (M. Ouaguenouni), meltem.ozturk@lamsade.dauphine.fr (M. Öztürk), olivier.spanjaard@lip6.fr (O. Spanjaard).

¹ A significant part of the work presented here has been carried out while Meltem Öztürk was on delegation at LIP6.

but solely on the stated preferences. Hence, we design an *ordinally robust* elicitation 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) and robust ordinal regression (Greco, Mousseau, & Słowiński, 2008)) by using a polyhedron of possible values for the parameters, reflecting the uncertainty about them. As a consequence, 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 wrong prediction, although a compromise must obviously be made between the reliability of the prediction and the predictive power of the learned model).

Elicitation setting. The input of our elicitation procedure is a learning set consisting of pairwise comparisons of various alternatives. More precisely, we consider an offline elicitation setting (passive learning) where we assume that a dataset of comparison examples is available, from which the parameters of the preference model are (partially) specified. This is a separate framework from the online elicitation setting (active learning) where we would incrementally select pairwise preference queries to enrich the learning set. The output of the elicitation procedure consists of pairwise comparisons that were not present in the learning set, which we call (preference) *predictions* hereafter. Note that, in some cases, the model may choose not to provide a prediction. The elicitation procedure thus results in a strict partial order on the alternatives.

Organization of the paper. After an overview of the related work (Section 2), we present the θ -additive utility model (Section 3), as well as the ordinal dominance relation inferred from it, based on the knowledge of a collection of preference examples. We then show how to determine whether a subset dominates another subset given the known pairwise preferences of the DM (Section 4), which enables to make preference predictions. The paper ends with numerical tests on synthetic and real-world preference data, and comparison with other preference learning methods (Section 5).

2. Related work

Preference elicitation (see e.g. Dias, Morton, & Quigley, 2018) and preference learning (see e.g. Fürnkranz & Hüllermeier, 2003, Corrente, Greco, Kadziński, & Słowiński, 2013) have been studied for a long time in operations research and artificial intelligence. This is a prerequisite in many applications across a wide range of fields, such as recommender systems, banking, financial management, chemistry, energy resources, health, investments, and industrial location (Andreopoulou, Koliouka, & Zopounidis, 2017). Several issues can be tackled in preference elicitation, and we detail three of them in the remainder of this section.

2.1. Incremental preference elicitation

To alleviate the cognitive burden on the DM, an incremental elicitation method can be adopted, in which comparison examples are interactively generated with the DM, to determine a necessary “optimal” alternative.

A pioneering work in this matter is the ISMAUT method (*Imprecisely Specified MultiAttribute Utility Theory*) (White, Sage, & Dozono, 1984). The utility $u(x)$ of a multiattribute alternative $x = (x_1, \dots, x_n)$ is defined as $u(x) = \sum_{i=1}^n w_i v_i(x_i)$, where the parameters are a weight vector $w = (w_1, \dots, w_n)$ and a vector $v = (v_1, \dots, v_n)$ of local value functions. When the parameter values are unknown, the (weak) preference statements made by a DM makes it possible to circumvent the set of her possible utility functions to $U = \{u : u(x) \geq u(y)\}$. An alternative x is then said

to *dominate* an alternative y if $u(x) \geq u(y)$ for all $u \in U$, and there exists u for which $u(x) > u(y)$.

The idea of defining a polyhedron of uncertainty (corresponding to U in the ISMAUT method) on the parameters of a multiattribute utility function goes back at least to the work of Charnetski and Soland (1978), which states that $A > B$ 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 . Another way to infer preferences from the polyhedron of uncertainty is to consider that the preferred alternative is within the set of nondominated alternatives and to obtain enough preferential information for the nondominated set to be reduced to a singleton.

In an active learning setting, i.e., when the pairwise comparisons are chosen by the analyst by directly querying the DM, and beyond the specific case of additive multiattribute utility, it is possible to choose the preference queries so as to reduce the nondominated set to a singleton as quickly as possible. In this way, the cognitive burden is limited for the DM. Popular approaches in this concern are to ask queries that yield to approximately halve the size of U (Ghosh & Kalagnanam, 2003; Iyengar, Lee, & Campbell, 2001), to use the *current solution strategy* (e.g., Benabbou, Leroy, Lust, & Perny, 2021; Boutilier, Patrascu, Poupart, & Schuurmans, 2006; Wang & Boutilier, 2003), i.e., a strategy based on the minmax regret criterion for choosing the next query, or to use a heuristic strategy for prioritizing the queries (Ciomek, Kadziński, & Tervonen, 2017).

A large body of literature in operations research deals with the same idea of taking advantage of a dominance relation between alternatives when the utility function is imprecisely specified. Notably, in multicriteria decision making (where x is a vector of criterion values instead of attribute values), *Stochastic Multiobjective Acceptability Analysis* (often abbreviated by SMAA (Lahdelma, Hokkanen, & Salminen, 1998) and SMAA-2 (Lahdelma & Salminen, 2001)) take into consideration the whole set of parameter values compatible with the known pairwise preference comparisons, in order to provide an index measuring the acceptability of an alternative in function of its rank among the available alternatives according to the different parameter value scenarios.

Besides, a whole stream of research has emerged under the name of *Robust Ordinal Regression* (often abbreviated by ROR). The pioneering work of Greco et al. (2008) aimed to introduce robustness in the UTA method, proposed by Jacquet-Lagrèze and Siskos in a seminal paper (Jacquet-Lagrèze & Siskos, 1982). The UTA method is an ordinal regression method using linear programming to estimate the parameters of an additive utility function, by minimizing the amount of total deviation from the inequalities between utility values induced by the known pairwise comparisons. Among other extensions of the methods, Greco et al. proposed to take into account all additive utility functions compatible with the pairwise comparisons, to enforce the robustness of conclusions.

Generally speaking, the name *ordinal regression* encompasses methods that aim at learning the parameters of a decision model from known pairwise comparisons between alternatives. The qualifying adjective *robust* is added to indicate that all parameter values compatible with the available preferential information are taken into account for preferential inference. Some restrictions on possible preferences are generally enforced, such as the monotony of the criteria in multicriteria decision-making, but extensions not making that assumption have also been considered (Corrente & Tasiou, 2023; Ghaderi, Ruiz, & Agell, 2017). Additionally, note that the terminology of the ROR framework slightly differs from the one used by White et al. (1984). The dominance relation is indeed called *necessary weak preference relation*. Furthermore, a *possible weak preference relation* between alternatives is also introduced, which holds as soon as there exists a utility function $u \in U$ such that $u(x) \geq u(y)$ (x is possibly weakly preferred to y). The ROR approach aims to provide a robust recommendation in terms of necessary and possible

preference relations, and the method we present in the remainder of the paper is in line with this concern.

As mentioned earlier, our method extends the idea to take into account not only all the possible *parameter values* compatible with the preferential data, but also all the possible *parameter sets*. For more details on ROR, the interested reader may refer to the two surveys by Corrente et al. (2013), Corrente, Greco, Kadziński, and Słowiński (2014), or to the recent survey by Kadziński (2022).

2.2. Incorrect preference statements and inconsistencies

The cognitive load placed on the decision makers can lead them to “incorrect” preference statements, or statements that, taken together, are not compatible with the utility model under consideration. This requires the development of elicitation methods capable of taking account of these potential errors and inconsistencies.

Bayesian approaches have been considered in this matter (Bourdache, Perny, & Spanjaard, 2019; Guo & Sanner, 2010), as well as possibilistic approaches (Adam & Destercke, 2021; Mousseau, Figueira, Dias, da Silva, & Clmaco, 2003). The idea in these works is, based on known pairwise comparisons, to specify an uncertainty measure over the parameter space, instead of a polyhedron of parameter values compatible with the available preferential information. In the context of incremental elicitation of a multiattribute utility function, an intermediate approach has been advocated by Sauré and Vielma (2019), namely to update an ellipsoidal credibility region computed from a multivariate normal distribution over the space of parameters. They showed how to use mixed integer programming to determine queries that are likely to reduce the volume of the credibility region, to gain more insights on the preferences of the DM.

2.3. Beyond the additive utility model

An important topic in decision making and preference elicitation is the ability to cope with preferences that cannot be represented by using an additive utility function. We focus on this topic in this paper, in the specific setting of subset comparisons. More precisely, we study the elicitation of a set function taking into account positive and negative interactions between elements, intending to predict pairwise preferences between subsets.

The Choquet integral is the most studied decision model for taking into account positive and negative interactions between criteria in multicriteria decision making (Grabisch, 1996). It turns out that a Choquet integral defined on binary vectors representing subsets can be viewed as a set function. Note that a Choquet integral is parameterized by a capacity ν on the criteria set N , i.e., a set function on N that is *monotone* ($A \subseteq B \Rightarrow \nu(A) \leq \nu(B)$) and *normalized* ($\nu(N) = 1$). As will become clear in the remainder of the paper, we do not impose such constraints in the model we consider.

Angilella, Corrente, and Greco (2015) developed a SMAA methodology to make robust recommendations based on the Choquet integral, by exploring the space of capacities and common utility scales compatible with the DM’s preference information (the criteria values must be mapped to a common scale to be aggregated by a Choquet integral). This work has been extended to the hierarchical Choquet integral preference model (Angilella, Corrente, Greco, & Słowiński, 2016), with a methodology combining SMAA and ROR. Let us mention some other recent works dealing with the elicitation of the parameters of a Choquet-related aggregation function: Bresson, Cohen, Hüllermeier, Labreuche, and Sebag (2020) use a perceptron approach to learn the parameters of a 2-additive hierarchical Choquet integral, while Herin, Perny, and Sokolovska (2023) propose an algorithm to learn sparse Möbius representations from preference examples, without a prior k -additivity assumption.

In the rest of the paper, given a “static” training set of examples of pairwise DM’s preferences, we propose a ROR method that circumvents the possible parameter sets and the possible parameter values,

which yields a dominance relation between subsets (corresponding to a necessary strict preference relation in the ROR terminology). This dominance relation is induced by the simplest parameter sets (to be formalized later). We predict pairwise comparisons between alternatives from this dominance relation. Note that predicting a comparison between alternatives can be framed as a binary classification problem by considering, 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. In this setting, many approaches have been proposed by the preference learning community in AI, going from perceptrons (Dragone, Teso, & Passerini, 2017) to Gaussian processes (Chu & Ghahramani, 2005) or Support Vector Machines (SVM) (Domshlak & Joachims, 2005).

An important feature of our elicitation procedure is that it may lead to not making predictions for some pairwise comparisons if the available preferential information is not conclusive enough. Other classification models may not predict a class for some examples, either because of an ambiguity in the class to predict (ambiguity rejection) or because the example is too far from the examples that are in the learning set (novelty rejection). These approaches have been grouped under the umbrella of *learning with rejection*, and are generally used in safety-sensitive domains, e.g. to predict a disease in medical applications (Kompa, Snoek, & Beam, 2021). For a complete review of learning with reject option, we refer the reader to the survey made by Hendrickx, Perini, Van der Plas, Meert, and Davis (2024).

To our knowledge, the works closest to our own are those of Domshlak and Joachims (2005) and Bigot, Fargier, Mengin, and Zanuttini (2012). Similarly to the preference model we consider, Domshlak and Joachims study a function that could represent any weak order on the alternatives. More precisely, they consider a multiattribute utility function that is a sum of 4^n subutilities over subsets of attribute values, where n is the number of attributes. The subutility values are then learned using an efficient SVM approach based on the *kernel trick* (see e.g., Schölkopf & Smola, 2002). Bigot et al. study the use of generalized additively independent decompositions of utility functions (Fishburn, 1970; Gonzales & Perny, 2005). They give a PAC-learner² that is polynomial time if a constant bound is known on the degree of the function, where the *degree* is the size of the greatest subset of attributes in the decomposition. Yet, both works do not fit the robust ordinal learning framework we consider in this work.

3. From the θ -additive model to \sqsubseteq -ordinal dominance

Given a set $\mathcal{F} = \{a_1, a_2, \dots, a_n\}$ of elements, we aim to reason on the preferences of the DM on a set \mathcal{A} of subsets $A \subseteq \mathcal{F}$, representing alternatives. The characteristic vector \bar{A} of a subset A is the n -dimensional binary vector whose i^{th} component is 1 if $a_i \in A$, and 0 otherwise. For instance, the characteristic vector of $A = \{a_1, a_2, a_4\}$ is $\bar{A} = (1, 1, 0, 1)$ if $\mathcal{F} = \{a_1, a_2, a_3, a_4\}$. In the following, we may use one or the other notation for describing a subset. Here are some examples of alternatives represented by subsets:

- If \mathcal{F} is a set of reference users expressing opinions on cultural products (e.g., movies), a cultural product may be represented by the subset A of reference users in \mathcal{F} that have a positive opinion on it, i.e., $a_i \in A$ if reference user a_i has a positive opinion on it, otherwise $a_i \notin A$.
- If \mathcal{F} is the set of players in a squad, a team lineup may be represented by the subset A of players that compound it.
- If \mathcal{F} is a set of binary features of technological products (e.g., smartphones), a technological product may be represented by a subset A of features, i.e., $a_i \in A$ if the product has feature a_i , otherwise $a_i \notin A$.

² PAC stands for Probably Approximately Correct.

We assume for simplicity that there are no two distinct alternatives corresponding to the same subset $A \subseteq \mathcal{F}$, which implies in particular that $2^{|\mathcal{F}|} \geq |\mathcal{A}|$. We infer strict pairwise preferences from strict preferences given by a DM on some subset of alternatives in \mathcal{A} , and we use this training set of pairwise preferences on alternatives (each viewed as a subset) to elicit the parameters of a utility function f defined on \mathcal{A} . The role of the utility function f is to represent the (unknown) strict weak order on \mathcal{A} corresponding to the DM's preferences, with $A > B$ iff $f(A) > f(B)$ and $A \sim B$ iff $f(A) = f(B)$. The relationship $A \sim B$, which occurs if neither $A > B$ nor $B > A$, should be interpreted here as incomparability and not as indifference.

We do not perform a full elicitation of the parameters of f , but we consider an uncertainty set of parameters values consistent with the known preferences of the DM, as in robust ordinal regression. If $f(A) > f(B)$ for all parameters values in this uncertainty set, then A is predicted to be strictly preferred to B . Actually, we do not only learn the parameters values, but also the components of the parameter set themselves, as we explain below.

3.1. The θ -additive model

Before coming to the proposed θ -additive model, we first recall the standard additive utility model, and its extension, the k -additive utility model.

The additive and k -additive utility models. As the DM's preferences over \mathcal{A} are modeled as a strict weak order, there exists a real-valued function f such that $\forall A, B \in \mathcal{A}, f(A) > f(B) \Leftrightarrow A > B$. Many models assume that f can be represented in a compact way using some sort of additivity property. The simplest and most used one is the additive model (Fishburn, 1970). This model makes the strong assumption that we can find a parameter value $v(a) \in \mathbb{R}$ for each element $a \in \mathcal{F}$ such that for all $A \in \mathcal{A}$, the utility of A is $f(A) = \sum_{a \in A} v(a)$. This assumption is strong because it implies that there is no interaction between the elements. A weaker assumption is that of k -additivity where we suppose the existence of a parameter $v(S) \in \mathbb{R}$ 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 $A \in \mathcal{A}$, $f(A) = \sum_{S \in [\mathcal{F}]^k} I_A(S)v_S$, where $I_A(S) = 1$ if $S \subseteq A$ and 0 otherwise, and v_S is an abbreviation for $v(S)$. Obviously, the 1-additive model amounts to the additive model. Taking k strictly greater than 1 makes it possible to account for (positive or negative) synergies between subsets of k or less elements. For example, the 2-additive model makes it possible to account for binary synergies. The utility of the alternative $A = (1, 1, 0, 1)$ with the 2-additive model is $f(A) = v(\{a_1\}) + v(\{a_2\}) + v(\{a_4\}) + v(\{a_1, a_2\}) + v(\{a_1, a_4\}) + v(\{a_2, a_4\})$. If there is a positive synergy between a_1 and a_2 then $f(\{a_1, a_2\}) > v(\{a_1\}) + v(\{a_2\})$ holds because $f(\{a_1, a_2\}) = v(\{a_1\}) + v(\{a_2\}) + v(\{a_1, a_2\})$. Note incidentally that $f(\{a_1, a_2\}) \neq v(\{a_1, a_2\})$. 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, Marichal, & Roubens, 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. Given a set $\theta \subseteq 2^{\mathcal{F}}$, and a set function $v : \theta \rightarrow \mathbb{R}$, we assume that f is of the form $f(A) = \sum_{S \in \theta} I_A(S)v_S$, where v_S stands again for $v(S)$. We call this the θ -additive model. For this model, we may also use the notation $f_{\theta,v}(A)$ instead of $f(A)$. The 1-additive (resp. k -additive) model is the special case in which $\theta = [\mathcal{F}]^1$ (resp. $\theta = [\mathcal{F}]^k$).

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

$$\begin{aligned} & \{a_2, a_3, a_4\} > \{a_1, a_3, a_4\} > \{a_1, a_2, a_4\} > \{a_3, a_4\} > \{a_2, a_4\} \\ & > \{a_2, a_3\} > \{a_1, a_4\} > \{a_1, a_3\} > \{a_1, a_2\} > \{a_4\} > \{a_3\} > \{a_2\} \\ & > \{a_1\} > A = \{a_1, a_2, a_3, a_4\} \sim \emptyset > B = \{a_1, a_2, a_3\}. \end{aligned}$$

Table 1

For each (n, k) , the table shows the proportion of linear extensions of the inclusion order on the power set of n elements that are compatible with a k -additive model.

Degree	$k = 1$	$k = 2$	$k = 3$
$n = 3$	25%	100%	100%
$n = 4$	0.02%	64.19%	100%

These preferences can be explained by a clear negative synergy when a_1, a_2 , and a_3 are chosen together (in A and B). Interestingly, instead of using a complete 3-additive model, which would require the definition of 14 parameters, this strict weak order can be obtained by using a θ -additive model with $\theta = \{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}, \{a_1, a_2, a_3\}\}$ and $v_{\{a_1\}} = 1, v_{\{a_2\}} = 2, v_{\{a_3\}} = 3, v_{\{a_4\}} = 4, v_{\{a_1, a_2, a_3\}} = -10$. This allows us to benefit from the expressiveness offered by 3-additivity while restricting the number of parameters.

Remark 1. Fishburn and Lavalley (1996) have pointed out that the 2-additive model offers a powerful corrective to the 1-additive model for comparing subsets, even if there are higher-order synergies. It seems to us, however, that this statement needs to be tempered. Many linear orderings of a power set can indeed not be represented by a 2-additive model. Even if we only consider the linear extensions of the inclusion order on the power set of $\{1, 2, 3, 4\}$ (the compatibility of preferences between sets with the inclusion relation seems a reasonable assumption), about one-third of the linear extensions require ternary interactions to be taken into account to obtain a compatible model. Put another way, only two-thirds of the linear extensions are compatible with a model of degree 3 for $n = 4$ elements. Furthermore, very few linear extensions (0.02%) can be represented using a purely additive model (without synergies). Table 1 synthesizes the results obtained for $n = 3$ and $n = 4$ (we did not go beyond $n = 4$ in our numerical tests because there already are about 1.5×10^{19} linear extensions to consider for $n = 5$). Note that relaxing compatibility with the inclusion relation would further increase the necessary degree of interactions to represent the permutations of the power set of $\{1, \dots, n\}$.

3.2. The θ -ordinal dominance relation

In our elicitation setting, we assume that we have only access to a partial set R of strict pairwise preferences provided by the DM.³ This set may contain only a few comparisons. Our aim is to use these observed preferences 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 $(A, B) \in R \Leftrightarrow A > B$.

Moreover, given θ , the set of value functions on θ that are compatible with the preferences observed in R is denoted by V_{θ}^R :

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

Note that, for a given θ , this set V_{θ}^R can be either empty or composed of an infinity of possible value functions on θ . Notably, if this set is empty then the preferences of the user cannot be represented by a θ -additive function. We denote by Θ^R the set $\{\theta \mid V_{\theta}^R \neq \emptyset\}$, i.e., the θ 's such that the preferences in R are consistent with a θ -additive function.

Unfortunately, given $\theta \in \Theta^R$ such that $V_{\theta}^R \neq \emptyset$, a pair $\{v, v'\}$ of value functions in V_{θ}^R may lead to infer opposite preferences, as illustrated below.

³ While we restrict our analysis to a set R compounded of strict pairwise preferences, our setting could accommodate negative preferences of the type $\neg(A > B)$ with few modifications. The preference statement $\neg(A > B)$ would yield the constraint $f_{\theta,v}(B) \geq f_{\theta,v}(A)$.

Example 2. 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 $\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}$:

$\{a_4\} > \{a_3\} > \{a_2\} > \{a_1\}$, or equivalently:

$$R = \{(\{a_4\}, \{a_3\}), (\{a_4\}, \{a_2\}), (\{a_4\}, \{a_1\}), (\{a_3\}, \{a_2\}), (\{a_3\}, \{a_1\}), (\{a_2\}, \{a_1\})\}.$$

The two additive functions v and v' defined by:

$$v(\{a_1\}) = 1, v(\{a_2\}) = 2, v(\{a_3\}) = 3, v(\{a_4\}) = 5$$

$$\text{and } v'(\{a_1\}) = 1, v'(\{a_2\}) = 3, v'(\{a_3\}) = 4, v'(\{a_4\}) = 5$$

are both in V_θ^R , but we infer $\{a_1, a_4\} > \{a_2, a_3\}$ from v while we infer $\{a_2, a_3\} > \{a_1, a_4\}$ from v' .

This example shows that, given R , choosing a specific function $v \in V_\theta^R$ can lead to infer preferences that are only related to this arbitrary choice (Bartee, 1971). Our aim is to infer preferences for pairs outside R in a reliable way by eliminating such arbitrary choices. In this purpose, we turn to a robust ordinal regression approach based on the observed preferences in R .

Fishburn and Lavalley (1996) showed how one can obtain an *ordinal dominance relation* from a partially specified 2-additive numerical model. We now explain how their idea can be extended to a θ -additive model.

Definition 1. Let \mathcal{F} be a set of elements, $\mathcal{A} \subseteq 2^{\mathcal{F}}$ a set of subsets and R a set of pairs $(A, B) \in \mathcal{A}^2$ where $(A, B) \in R \Leftrightarrow A > B$. Given $\theta \in \Theta^R$, the θ -ordinal dominance relation, denoted by $>_\theta^R$, is defined for $A, B \in \mathcal{A}$ by:

$$A >_\theta^R B \Leftrightarrow \forall v \in V_\theta^R, f_{\theta,v}(A) > f_{\theta,v}(B).$$

The θ -ordinal dominance relation is independent from the choice of a specific $v \in V_\theta^R$. Naturally, $(A, B) \in R \Rightarrow A >_\theta^R B$. Nevertheless, note that the binary relation $>_\theta^R$ is obviously partial, and we define the incomparability relation \sim_θ^R as:

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

If $A >_\theta^R B$ then we 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

We conclude this section by mentioning some properties of $>_\theta^R$:

- Unlike $>$, the relation $>_\theta^R$ is not a strict weak order: it is asymmetric but it may not be complete nor negatively transitive. The absence of preference prediction may occur in two not equivalent situations: either A and B belong to the same incomparability class of the (unknown) strict weak order $>$ on \mathcal{A} , i.e., $A \sim B$, or there is not enough preferential information in R to conclude that $A > B$ or $B > A$.
- Since the θ -ordinal dominance relation depends on the preference set R and on the model θ , the relation $>_\theta^R$ evolves when θ or R are restricted or extended. In particular, if $\theta' \subseteq \theta$ then any prediction that is yielded using θ -ordinal dominance is also yielded using θ' -ordinal dominance; thus, if $V_\theta^R \neq \emptyset$ and $V_{\theta'}^R \neq \emptyset$, then θ' appears as more appealing from a preference learning standpoint since it allows more predictions to be made. Furthermore, one could prefer θ' over θ because of the philosophical principle of parsimony (e.g. Blumer, Ehrenfeucht, Haussler, & Warmuth, 1987).

A more formal and detailed description of the properties of $>_\theta^R$ can be found in the supplementary material (see Appendix).

3.3. The \sqsubseteq -ordinal dominance relation

The θ -ordinal dominance relation depends on the choice of a specific set $\theta \in \Theta^R$ but, as shown in the following example, several θ 's may be in Θ^R .

Example 3. Assume that R consists of all pairwise preferences resulting from $>$ in **Example 1**. Setting $\theta = \{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}\}$ yields then $V_\theta^R = \emptyset$. In contrast, setting $\theta_1 = \{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}, \{a_1, a_2, a_3\}\}$ yields $V_{\theta_1}^R \neq \emptyset$. Actually, there are many other sets θ compatible with the preferences in R : it can be shown⁴ that $\Theta^R = \{\theta : \theta_1 \subseteq \theta\}$ for this example.

The question that naturally arises is whether we could find two different models $\theta_1, \theta_2 \in \Theta^R$ that are both compatible with the observed preferences in R and such that $A >_{\theta_1}^R B$ and $B >_{\theta_2}^R A$ for a pair of alternatives $(A, B) \in \mathcal{A}^2$. Unfortunately, this situation may indeed happen:

Example 4. Let $R = \{(\{a_1\}, \{a_2\})\}$, $\theta_1 = \{\{a_1\}\}$ and $\theta_2 = \{\{a_2\}\}$. Note that both θ_1 and θ_2 belong to Θ^R . If we consider $\theta_1 = \{\{a_1\}\}$, the set $V_{\theta_1}^R$ is compounded of value functions v defined on θ_1 such that $v(\{a_1\}) > 0$. Hence, for all $v \in V_{\theta_1}^R$ we have $f_{\theta_1,v}(\{a_1, a_2\}) = v(\{a_1\}) > 0 = f_{\theta_1,v}(\emptyset)$ and thus $\{a_1, a_2\} >_{\theta_1}^R \emptyset$. Conversely, if we consider $\theta_2 = \{\{a_2\}\}$, the set $V_{\theta_2}^R$ is compounded of value functions v defined on θ_2 such that $v(\{a_2\}) < 0$. This yields $f_{\theta_2,v}(\{a_1, a_2\}) = v(\{a_2\}) < 0$ for each $v \in V_{\theta_2}^R$ and thus $\emptyset >_{\theta_2}^R \{a_1, a_2\}$.

In what follows, we define a new dominance relation taking into account the plurality of models compatible with the observed preferences.

Note that there always exists a θ able to represent R (at worst, $\theta = 2^{\mathcal{F}}$) and that if a θ -additive model is compatible with R , then any θ' -additive model with $\theta \subseteq \theta'$ is also compatible with R . For this reason, the number of sets θ compatible with the observed preferences may be very large.

For this reason, we start by restricting the set of models to take into account. In this purpose, we need a binary relation \sqsubseteq on Θ^R , such that $\theta \sqsubseteq \theta'$ if θ is considered simpler than θ' . Our idea is to only consider sets θ that are minimal according to such a binary relation, i.e., θ such that $\nexists \theta' \in \Theta^R$ for which $\theta' \sqsubseteq \theta$. This is motivated by the philosophical principle of parsimony that the simpler of two explanations is to be preferred (Occam's razor (Blumer et al., 1987)). Different possible definitions for \sqsubseteq will be discussed upon in the following subsection.

We call \sqsubseteq -simplest θ of Θ^R the parameter sets $\theta \in \Theta^R$ which are minimal w.r.t. \sqsubseteq , and we denote by Θ_\sqsubseteq^R their set. Based on Θ_\sqsubseteq^R , we define the \sqsubseteq -ordinal dominance relation:

Definition 2. Let \mathcal{F} be a set of elements, $\mathcal{A} \subseteq 2^{\mathcal{F}}$ a set of subsets and R a set of pairs $(A, B) \in \mathcal{A}^2$ where $(A, B) \in R \Leftrightarrow A > B$. The \sqsubseteq -ordinal dominance relation, denoted by $>_\sqsubseteq^R$, is defined, for $A, B \in \mathcal{A}$, as follows:

$$A >_\sqsubseteq^R B \Leftrightarrow \forall \theta \in \Theta_\sqsubseteq^R, A >_\theta^R B,$$

$$\Leftrightarrow \forall \theta \in \Theta_\sqsubseteq^R, \forall v \in V_\theta^R, f_{\theta,v}(A) > f_{\theta,v}(B).$$

In other words, A \sqsubseteq -ordinally dominates B if A θ -ordinally dominates B according to all θ in Θ_\sqsubseteq^R , i.e., all the \sqsubseteq -simplest θ 's of Θ^R .

3.4. Different definitions for \sqsubseteq

We say that a relation \sqsubseteq is based on a function ξ when $\theta \sqsubseteq \theta'$ if and only if $\xi(\theta) \leq \xi(\theta')$. Several aspects can be taken into account to define ξ :

⁴ It has been computer tested by brute force enumeration.

• A first idea is to favor parameter sets θ that minimize the complexity of synergies between the attributes. To measure this complexity, we use the *degree* of θ , namely $\deg(\theta) = \max\{|S| : S \in \theta\}$ (i.e., the greatest cardinality of a subset of interacting attributes). This leads to the binary relation \sqsubseteq_{\deg} based on \deg , i.e., $\theta \sqsubseteq_{\deg} \theta' \Leftrightarrow \deg(\theta) \leq \deg(\theta')$.

• A second idea is to favor parameter sets θ having the *sparsest* possible representation (Zhang, Xu, Yang, Li, & Zhang, 2015), i.e., those which minimize $\text{card}(\theta) = |\theta|$. This choice yields the binary relation $\sqsubseteq_{\text{card}}$, which is the relation based on the function card , i.e., $\theta \sqsubseteq_{\text{card}} \theta' \Leftrightarrow \text{card}(\theta) \leq \text{card}(\theta')$.

• Alternatively, we define a binary relation combining the ideas of \sqsubseteq_{\deg} and $\sqsubseteq_{\text{card}}$ by considering both the number and the size of elements in a parameter set θ . In this purpose, we define \sqsubseteq_{ws} , the relation based on the function $\text{ws}(\theta) = \sum_{S \in \theta} |S|$, i.e., $\theta \sqsubseteq_{\text{ws}} \theta' \Leftrightarrow \text{ws}(\theta) \leq \text{ws}(\theta')$.

• Lastly, we define the binary relation \sqsubseteq_{lex} , defined by using lexicographically the binary relations \sqsubseteq_{\deg} , $\sqsubseteq_{\text{card}}$, and \sqsubseteq_{ws} , in this order. This relation could be seen as based on the function lex where $\text{lex}(\theta) = n4^n \deg(\theta) + n2^n \text{card}(\theta) + \text{ws}(\theta)$.

All those relations provide different refinements of the partial order induced by inclusion, i.e., $\theta \subseteq \theta' \Rightarrow \theta \sqsubseteq \theta'$ for $\sqsubseteq \in \{\sqsubseteq_{\deg}, \sqsubseteq_{\text{card}}, \sqsubseteq_{\text{ws}}, \sqsubseteq_{\text{lex}}\}$.

Example 5. Let $R = \{(\{a_1, a_2\}, \{a_3, a_4\}), (\{a_1, a_2\}, \{a_1, a_3\})\}$. It is easy to see that $V_\theta^R \neq \emptyset$ for $\theta = \{\{a_1, a_2\}\}$, which corresponds to a model of degree 2. However, we may prefer being consistent with a model of degree 1, even if there are more elements in it: $\theta' = \{\{a_1\}, \{a_2\}\}$ or $\theta'' = \{\{a_1\}, \{a_3\}\}$ or $\theta''' = \{\{a_2\}\}$. In this example, the minimal parameter set θ among $\theta', \theta'', \theta'''$ w.r.t. relation \sqsubseteq_{\deg} (resp. $\sqsubseteq_{\text{card}}, \sqsubseteq_{\text{ws}}, \sqsubseteq_{\text{lex}}$) is $\{\theta', \theta'', \theta'''\}$ (resp. $\{\theta'''\}$) in the three cases).

4. Preference prediction by using \sqsubseteq -ordinal dominance

Given a set R of pairwise preferences and a binary relation \sqsubseteq on Θ^R , the preference learning method we propose consists in predicting that a subset A is preferred to B if $A \succ_{\sqsubseteq}^R B$, i.e., A is preferred to B for all simplest models $\theta \in \Theta^R$ and value functions $v \in V_\theta^R$. The purpose of this section is to detail the procedure for determining whether $A \succ_{\sqsubseteq}^R B$. It is organized as follows:

• We show that determining if $A \succ_{\theta}^R B$ is polytime in $|R|$ and $|\theta|$, while determining if $A \succ_{\sqsubseteq}^R B$ amounts to testing whether $\Theta_{\sqsubseteq}^R \cap \Theta_{B \succ A}^R = \emptyset$, where $\Theta_{B \succ A}^R = \{\theta \in \Theta^R : B \succ_{\theta}^R A \text{ or } B \sim_{\theta}^R A\}$ (Section 4.1).

• As determining an explicit representation of Θ_{\sqsubseteq}^R is likely to be cumbersome (as the size of Θ_{\sqsubseteq}^R can be very large), we turn to an implicit representation based on the values $\deg(\theta)$, $\text{card}(\theta)$, $\text{ws}(\theta)$ for $\theta \in \Theta_{\sqsubseteq}^R$. We thus study the computational complexity of determining $\deg(\theta)$ (resp. $\text{card}(\theta)$, $\text{ws}(\theta)$, $\text{lex}(\theta)$) for $\theta \in \Theta_{\sqsubseteq}^R$ and $\sqsubseteq = \sqsubseteq_{\deg}$ (resp. $\sqsubseteq = \sqsubseteq_{\text{card}}, \sqsubseteq = \sqsubseteq_{\text{ws}}, \sqsubseteq = \sqsubseteq_{\text{lex}}$), showing that the former problem can be solved in polynomial time, while the others are NP-hard (Section 4.2).

• The implicit representation of Θ_{\sqsubseteq}^R is based on the following idea: if we know that $\theta_0 \in \Theta_{\sqsubseteq}^R$, then $\theta \in \Theta_{\sqsubseteq}^R \Leftrightarrow (\deg(\theta), \text{card}(\theta), \text{ws}(\theta)) = (\deg(\theta_0), \text{card}(\theta_0), \text{ws}(\theta_0))$. It is thus enough to determine a single model $\theta_0 \in \Theta_{\sqsubseteq}^R$ to be able to determine whether a model belongs to Θ_{\sqsubseteq}^R . This is why we propose a Mixed Integer Program (MIP) to compute a model $\theta \in \Theta_{\sqsubseteq}^R$, derived from a linear program for determining whether a model θ belongs to Θ^R (Section 4.3).

• We derive from it another MIP to compute a model in $\Theta_{\sqsubseteq}^R \cap \Theta_{B \succ A}^R$, concluding $A \not\succeq_{\sqsubseteq}^R B$ if it exists, $A \succ_{\sqsubseteq}^R B$ otherwise (Section 4.4).

Remark 2. For \sqsubseteq a refinement of \sqsubseteq_{\deg} and $\Theta_{\sqsubseteq}^R \subseteq 2^{[F]^2}$, the \sqsubseteq -ordinal dominance relation is more discriminant (i.e., yields at least as many pairwise preferences) than the ordinal dominance relation considered by Fishburn and Lavalley (1996), which corresponds to $[F]^2$ -ordinal dominance. The ordinal dominance relation indeed defines a parameter for each subset in $[F]^2$ while the \sqsubseteq -dominance relation considers the \sqsubseteq -simplest θ 's of Θ^R (i.e., all $\theta \in \Theta_{\sqsubseteq}^R$). As long as the preferences may be represented by a 2-additive model, the \sqsubseteq -simplest θ 's are all




Fig. 1. $(\Theta_{\sqsubseteq}^R \cap \Theta_{B \succ A}^R = \emptyset \Leftrightarrow A \succ_{\sqsubseteq}^R B)$ and $(\Theta_{\sqsubseteq}^R \cap \Theta_{A \succ B}^R = \emptyset \Leftrightarrow B \succ_{\sqsubseteq}^R A)$.

included in $[F]^2 = \{S \subseteq F : 1 \leq |S| \leq 2\}$. We recall that if $\theta' \subseteq \theta$, any prediction obtained by θ -ordinal dominance is also obtained by θ' -ordinal dominance. Hence, all preferences obtained by $[F]^2$ -ordinal dominance is also obtained by \sqsubseteq -ordinal dominance relation.

4.1. Determining whether $A \succ_{\theta}^R B$ and whether $A \succ_{\sqsubseteq}^R B$

We first show that, unsurprisingly, linear programming provides an operational tool for determining whether $A \succ_{\theta}^R B$. Viewing a value function on θ as a vector $v = (v_S)_{S \in \theta}$ where $v_S = v(S)$, the set V_θ^R corresponds to the polyhedron defined by the following linear constraints in the $|\theta|$ -dimensional parameter space (where each parameter v_S corresponds to a dimension)⁵:

$$\forall (X, Y) \in R, \sum_{S \in \theta} I_X(S)v_S - \sum_{S \in \theta} I_Y(S)v_S \geq 1.$$

For a given set R of strict pairwise preferences and a model $\theta \in \Theta^R$, checking whether $A \succ_{\theta}^R B$ can be evaluated in polynomial time in $|R|$ and in $|\theta|$ by solving the following linear program $\mathcal{P}_{A \succ_{\theta}^R B}$, where there is one variable $v_S \in \mathbb{R}$ for each pair $S \in \theta$:

$$(\mathcal{P}_{A \succ_{\theta}^R B}) \left\{ \begin{array}{l} \min \sum_{S \in \theta} I_A(S)v_S - \sum_{S \in \theta} I_B(S)v_S \\ \sum_{S \in \theta} (I_X(S) - I_Y(S))v_S \geq 1 \quad \forall (X, Y) \in R \setminus \{(A, B)\}, \\ v_S \in \mathbb{R} \quad \forall S \in \theta. \end{array} \right.$$

We have that $A \succ_{\theta}^R B$ if and only if the optimal value of $\mathcal{P}_{A \succ_{\theta}^R B}$ is strictly positive, as it implies that $\sum_{S \in \theta} I_A(S)v_S > \sum_{S \in \theta} I_B(S)v_S$ for all $v \in V_\theta^R$.

In contrast with this positive complexity result for θ -ordinal dominance, determining whether $A \succ_{\sqsubseteq}^R B$ by direct use of the definition of \sqsubseteq -ordinal dominance would require a high computational burden. We overcome this difficulty by reducing this problem to testing whether $\Theta_{\sqsubseteq}^R \cap \Theta_{B \succ A}^R$ is empty.

To achieve this reduction, let us study the relationships between $\Theta_{A \succ B}^R$, $\Theta_{B \succ A}^R$ and Θ_{\sqsubseteq}^R . For visual support, the reader may refer to Fig. 1. We recall that we denote by $\Theta_{B \succ A}^R$ the set $\{\theta \in \Theta^R : B \succ_{\theta}^R A \text{ or } B \sim_{\theta}^R A\}$. As one of the relations $A \succ_{\theta} B$ or $B \succ_{\theta} A$ or $A \sim_{\theta} B$ holds for any $\theta \in \Theta^R$, we have that $\Theta^R = \Theta_{A \succ B}^R \cup \Theta_{B \succ A}^R$. Consequently, $\Theta_{\sqsubseteq}^R \subseteq \Theta_{A \succ B}^R \cup \Theta_{B \succ A}^R$ because $\Theta_{\sqsubseteq}^R \subseteq \Theta^R$. Furthermore, $\Theta_{A \succ B}^R \cap \Theta_{B \succ A}^R = \{\theta \in \Theta^R : A \sim_{\theta} B\} \neq \emptyset$ as soon as there exists $\theta \in \Theta^R$ for which $A \sim_{\theta} B$.

To evaluate whether a \sqsubseteq -ordinal dominance relation holds between two subsets A and B , we examine if one of the following conditions holds: (i) $\Theta_{\sqsubseteq}^R \cap \Theta_{B \succ A}^R = \emptyset$, (ii) $\Theta_{\sqsubseteq}^R \cap \Theta_{A \succ B}^R = \emptyset$. We have indeed the following result:

Proposition 1. For any $A, B \subseteq F$, we have $A \succ_{\sqsubseteq}^R B \Leftrightarrow \Theta_{\sqsubseteq}^R \cap \Theta_{B \succ A}^R = \emptyset$.

⁵ The right hand side of the constraint is here set to 1, but it could be set to any strictly positive constant as utilities v_S are always compatible with R to within a positive multiplicative factor.

Proof. It follows from the following sequence of equivalences:

$$A \succ_{\sqsubseteq}^R B \Leftrightarrow \forall \theta \in \Theta_{\sqsubseteq}^R, A \succ_{\theta}^R B \Leftrightarrow \forall \theta \in \Theta_{\sqsubseteq}^R, B \not\succeq_{\theta}^R A \text{ and } A \sim_{\theta}^R B \\ \Leftrightarrow \Theta_{\sqsubseteq}^R \cap \Theta_{B \not\succeq_{\theta}^R A} = \emptyset. \quad \square$$

Symmetrically, we have obviously that $B \succ_{\theta}^R A \Leftrightarrow \Theta_{\sqsubseteq}^R \cap \Theta_{A \not\succeq_{\theta}^R B} = \emptyset$. To test whether $\Theta_{\sqsubseteq}^R \cap \Theta_{B \not\succeq_{\theta}^R A} = \emptyset$, the mathematical programming approach we propose applies to cases where relation \sqsubseteq is based on a function ξ . The approach starts by computing a *single* model $\theta \in \Theta^R$ minimizing $\xi(\theta)$, which is enough for determining the value $\xi(\theta)$ of *any* $\theta \in \Theta_{\sqsubseteq}^R$, as they all share the same optimal value $\xi(\theta)$. We now study the complexity of computing such an optimal θ in Θ^R . More precisely, we study the complexity of the following decision problem MIN- θ - ξ , for $\xi \in \{\text{card, ws, deg, lex}\}$ (as is well-known, the optimization problem is at least as hard as its decision variant):

MIN- θ - ξ

INPUT: A set \mathcal{A} of alternatives, a set $R = \{(A, B), A, B \in \mathcal{A}\}$ of strict pairwise preferences, an integer $\tau \in \mathbb{Z}^+$.

QUESTION: Does there exist $\theta \in \Theta^R$ such that $\xi(\theta) \leq \tau$?

4.2. Computational complexity of MIN- θ - ξ for $\xi \in \{\text{card, ws, lex, deg}\}$

We show here that MIN- θ - ξ is NP-hard for $\sqsubseteq \in \{\text{ws, card, lex}\}$, while it can be solved in polynomial time for $\sqsubseteq = \text{deg}$.

Theorem 1. MIN- θ -card and MIN- θ -ws are NP-complete.

Proof. The membership of MIN- θ -card to NP follows from the fact that $\min_{\theta} \text{card}(\theta) \leq 2|R|$ and checking that $\theta \in \Theta^R$ can be done in polynomial time in $|R|$ and $|\theta|$. Indeed, the parameter set $\theta = \{A \in \mathcal{A} : (A, \cdot) \in R \text{ or } (\cdot, A) \in R\}$ obviously belongs to Θ^R , and $|\theta| \leq 2|R|$. The proof that MIN- θ -ws belongs to NP is similar, based on the fact that $\min_{\theta} \text{ws}(\theta) \leq 2|R| \times n$.

To prove the NP-hardness, we use a reduction from Hitting Set:

HITTING SET

INPUT: Given a set of n elements: $\mathcal{X} = \{x_i\}_{1 \leq i \leq n}$, a family of m sets $S = \{S_i : S_i \subseteq \mathcal{X}, 1 \leq i \leq m\}$, and an integer $\tau \in \mathbb{Z}^+$.

QUESTION: Does there exist $\mathcal{X}' \subseteq \mathcal{X}$ such that $\forall S_i \in S, S_i \cap \mathcal{X}' \neq \emptyset$ and $|\mathcal{X}'| \leq \tau$?

Given an instance (\mathcal{X}, S, τ) of the Hitting Set problem, we define the following instance (\mathcal{A}, R, τ') of MIN- θ -card (resp. MIN- θ -ws).

We let $\mathcal{A} = S \cup \{\emptyset\}$, $\tau' = \tau$, and consider the following set of preferences:

$$R = \{(S, \emptyset) : S \in S\}.$$

Now we show that (\mathcal{X}, S, τ) is a yes-instance of Hitting Set iff (\mathcal{A}, R, τ') is a yes-instance of MIN- θ -card (resp. MIN- θ -ws). Note that a set θ belongs to Θ^R if and only if it satisfies the following condition:

$$\forall (S, \emptyset) \in R, \exists T \in \theta \text{ such that } T \subseteq S. \quad (C)$$

Indeed, each preferences in R can then be satisfied by assigning positive values to parameters entailed by the elements of θ . Moreover, note that if a set θ satisfies C and $\exists T \in \theta$ such that $|T| > 1$, then the set θ' obtained from θ by replacing T by any singleton $\{x\} \subset T$ also satisfies C. Hence, within the sets satisfying C and minimizing card, there exists a set θ' compounded only of singletons, minimizing both card and ws (because $\text{card}(\theta) = \text{ws}(\theta)$ if θ is compounded only of singletons). By taking $\mathcal{X}' = \{x : \{x\} \in \theta'\}$, we obtain a hitting set of size $|\mathcal{X}'| \leq \tau$. This yields the following conclusion: there exists a hitting set of size $s \leq \tau$ if and only if there exists a set θ satisfying C such that $\text{card}(\theta) = s$ (resp. $\text{ws}(\theta) = s$). This argument completes the proof. \square

The following result is a direct consequence of the previous one:

Corollary 1. MIN- θ -lex is NP-hard.

Proof. Given an instance (\mathcal{A}, R, τ) of the MIN- θ -card problem, we could solve for each degree $d \in \{0, 1, \dots, |\mathcal{F}|\}$ an instance (\mathcal{A}, R, τ') of the MIN- θ -lex problem where $\tau' = dn4^n + (\tau + 1)n2^n$. \square

In contrast, we show a polynomial-time complexity result for MIN- θ -deg, by resorting to the *kernel trick*, widely used in machine learning (see e.g., Schölkopf & Smola, 2002). Given a vector space \mathcal{X} of dimension $n_{\mathcal{X}}$ and a transformation function $\varphi : \mathcal{X} \rightarrow \mathcal{Y}$, where the dimension $n_{\mathcal{Y}}$ of vector space \mathcal{Y} is exponential in $n_{\mathcal{X}}$, the kernel trick consists in computing the scalar products $\langle \varphi(x), \varphi(y) \rangle$ of $x, y \in \mathcal{X}$ in polynomial time in $n_{\mathcal{X}}$, by using a kernel function $K(x, y)$ that returns the value $\langle \varphi(x), \varphi(y) \rangle$ without requiring to explicit $\varphi(x)$ and $\varphi(y)$. In our setting, \mathcal{X} is the set of characteristic vectors of subsets A of \mathcal{F} , and \mathcal{Y} the set of “augmented” characteristic vectors containing additional dimensions corresponding to binary values $I_A(S)$ for $S \in [\mathcal{F}]^{\tau}$ (more details in the proof). The complexity result is formulated as follows:

Theorem 2. MIN- θ -deg can be solved in polynomial time in $|R|$ and n .

Proof. Let (\mathcal{A}, R, τ) be an instance of MIN- θ -deg. We wish to determine if preferences in R can be represented by a θ -additive model with $\theta = [\mathcal{F}]^{\tau}$. For notational convenience, we set $\theta^{(\tau)} = [\mathcal{F}]^{\tau}$ and $n_{\tau} = |\theta^{(\tau)}| = \sum_{i=1}^{\tau} \binom{n}{i}$. We associate to $\theta^{(\tau)}$ the vector $\overline{\theta^{(\tau)}} = (S_1, \dots, S_{n_{\tau}})$, where subsets $S = \{a_{i_1}, \dots, a_{i_k}\}$ ($i_1 < \dots < i_k$) are indexed in lexicographic order of vectors $(|S|, i_1, \dots, i_k)$. For instance, if $\mathcal{F} = \{a_1, a_2, a_3\}$ and $\theta = \theta^{(3)}$ then $\overline{\theta} = (\{a_1\}, \{a_2\}, \{a_3\}, \{a_1, a_2\}, \{a_1, a_3\}, \{a_2, a_3\}, \{a_1, a_2, a_3\})$. Additionally, for a value function $v : \theta \rightarrow \mathbb{R}$, we denote by $\overline{v} = (v_{S_1}, \dots, v_{S_{n_{\tau}}})$ the vector of values associated to the elements of $\overline{\theta^{(\tau)}}$ ordered in the same fashion. Finally, given $A \in \mathcal{A}$, we denote by \overline{A}_{τ} the binary vector $\overline{A}_{\tau} = (I_A(S_1), \dots, I_A(S_{n_{\tau}}))$, where $I_A(S_i)$ is the indicator function of $S_i \in \theta^{(\tau)}$.

Problem MIN- θ -deg evaluates if the following proposition holds:

$$\exists \overline{v} \in \mathbb{R}^{n_{\tau}} \text{ s.t. } \forall (A, B) \in R; \overline{A}_{\tau} \overline{v}^T > \overline{B}_{\tau} \overline{v}^T.$$

A value vector \overline{v} of minimum norm can be determined by solving the following convex quadratic program:

$$\min_{\overline{v} \in \mathbb{R}^{n_{\tau}}} \frac{1}{2} \overline{v} \overline{v}^T \\ \text{s.t. } \overline{A}_{\tau} \overline{v}^T \geq \overline{B}_{\tau} \overline{v}^T + 1 \quad \forall (A, B) \in R$$

Using the same trick as Domshlak and Joachims (Domshlak & Joachims, 2005), instead of solving this program whose number n_{τ} of variables is not polynomial in the size of our instance of MIN- θ -deg (because τ is an input variable and not a constant), we consider its Wolfe dual defined by:

$$\max_{\alpha \in \mathbb{R}^{|R|}} \sum_{(A, B) \in R} \alpha_{(A, B)} - \frac{1}{2} \sum_{(A, B) \in R} \sum_{(C, D) \in R} \alpha_{(A, B)} \alpha_{(C, D)} (\overline{A}_{\tau} - \overline{B}_{\tau})(\overline{C}_{\tau} - \overline{D}_{\tau})^T \\ \text{s.t. } \alpha \geq 0$$

By defining the kernel function $K^{(\tau)}(A, B) = \overline{A}_{\tau} \overline{B}_{\tau}^T$, the previous program can be written as:

$$\max_{\alpha \in (\mathbb{R}^+)^{|R|}} \sum_{(A, B) \in R} \alpha_{(A, B)} - \frac{1}{2} \sum_{(A, B) \in R} \sum_{(C, D) \in R} \alpha_{(A, B)} \alpha_{(C, D)} \\ (K^{(\tau)}(A, C) - K^{(\tau)}(A, D) - K^{(\tau)}(B, C) + K^{(\tau)}(B, D))$$

which can be solved in polynomial time in $|R|$ and n provided that $K^{(\tau)}(X, Y)$ can be evaluated in polynomial time in n without expliciting X and Y .

Indeed, since the reformulation yields a convex quadratic program of polynomial size in the input data, the problem can then be solved in polynomial time (by polynomial time solvability of convex quadratic

programming (Kozlov, Tarasov, & Khachiyan, 1979, 1980)). We now prove that $K^{(\tau)}(X, Y)$ can be efficiently computed without expliciting X and Y . Let k be the size of the intersection between X and Y , i.e., $k = |X \cap Y|$. Note that $K^{(\tau)}(X, Y)$ counts the number of parameters of $\theta^{(\tau)}$ that are subsets of both X and Y . We conclude by noting that the number of such elements corresponds to $\sum_{i=1}^{\tau} \binom{k}{i}$, i.e., the number ($< 2^k$) of non-empty subsets of size less than or equal to τ in $X \cap Y$. \square

Remark 3. Note that Tehrani, Strickert, and Hüllermeier (Tehrani et al., 2014) and Herin et al. (Herin et al., 2023) have proposed kernel functions $K(x, y)$ that return the scalar product $\langle \varphi(x), \varphi(y) \rangle$ of augmented vectors $\varphi(x), \varphi(y)$ used to obtain an additive expression $\langle m, \varphi(x) \rangle$ of a discrete Choquet integral $C(x)$, where m is the vector of Möbius masses obtained from the capacity used in $C(x)$. It turns out that there is a close link between $f_{\theta, v}$ and a Choquet integral $C(x)$ expressed as $\langle m, \varphi(x) \rangle$ (note however that we do not impose the constraints on the $v(S)$ values ensuring the monotonicity of the capacity, or the normalization constraint $\sum_S v(S) = 1$). However, their kernel functions do not use the same calculations as ours: we take advantage of the particular case we study, where all components of x take binary values, to compute the kernel function in $O(n)$ instead of $O(n^2)$.

Algorithm 1 takes as input a set R of strict pairwise preferences and computes $\min\{\text{deg}(\theta) : \theta \in \Theta^R\}$ by solving a sequence of convex quadratic programs establishing whether there exists $\theta \in \Theta^R$ such that $\xi(\theta) = \tau$ (which holds if the optimal value of the program is bounded). The variable τ is gradually incremented from 1. At each iteration, the objective function parameters are updated by using the kernel trick, which makes the procedure polynomial-time in $|R|$ and n . More precisely, the while loop makes at most n calls to a convex quadratic programming solver (which operates in polynomial time in the size of the program (Kozlov et al., 1979, 1980)), and each iteration of the while loop requires to update $|R|^2$ components $Q[A, B, C, D]$. To perform efficiently the updates, we can precompute the binomial coefficients in time $O(n^3)$, by using dynamic programming to determine the Pascal triangle. The Pascal triangle indeed involves n^2 binomial coefficients, each one computed by an addition in $O(n)$ (as the values of the coefficients are upper bounded by 2^n , thus encoded on n bits). The overall complexity of Algorithm 1 is dominated by the $|n|$ solutions of convex quadratic programs.

Algorithm 1 Compute $\min\{\text{deg}(\theta) : \theta \in \Theta^R\}$

Input: set R of strict pairwise preferences

Output: $\min\{\text{deg}(\theta) : \theta \in \Theta^R\}$

```

τ ← 1
for (A, B) ∈ R do
    for (C, D) ∈ R do
        for (A, B, C, D) do
            Q[A, B, C, D] ← |A ∩ C| − |A ∩ D| − |B ∩ C| + |B ∩ D|
        end for
    end for
while max_{α ≥ 0} ∑_{(A,B) ∈ R} α_{(A,B)} − 1/2 ∑_{(A,B) ∈ R} ∑_{(C,D) ∈ R} α_{(A,B)} α_{(C,D)} Q[A, B, C, D] is unbounded do
    ▷ the α_{(X,Y)}'s are the variables of the convex quadratic program
    ▷ α ≥ 0 means that α_{(X,Y)} ≥ 0 for all (X, Y) ∈ R
    ▷ Q contains the coefficients of the objective function, updated at each iteration
    τ ← τ + 1
    for (A, B) ∈ R do
        for (C, D) ∈ R do
            Q[A, B, C, D] ← Q[A, B, C, D] +  $\binom{|A \cap C|}{\tau} - \binom{|A \cap D|}{\tau} - \binom{|B \cap C|}{\tau} + \binom{|B \cap D|}{\tau}$ 
        end for
    end for
return τ

```

4.3. Computing $(\text{deg}(\theta), \text{card}(\theta), \text{ws}(\theta))$ for $\theta \in \Theta_{\text{lex}}^R$

As all models $\theta \in \Theta_{\text{lex}}^R$ share the same vector $(\text{deg}(\theta), \text{card}(\theta), \text{ws}(\theta))$, it is enough to compute a single model $\theta \in \Theta_{\text{lex}}^R$ to deduce this vector, which will be required to determine whether $A \succ_{\text{lex}}^R B$. The negative complexity result (Corollary 1) regarding the computation of a model

$\theta \in \Theta_{\text{lex}}^R$ does not prevent us from proposing an exact solution method that will prove efficient in practice. For this purpose, we first present a Linear Program (LP) allowing us to determine in polynomial time in $|R|$ and $|\theta|$ whether $\theta \in \Theta^R$, given a model θ and a set R of strict pairwise preferences. From this LP, we will then develop a MIP formulation for computing $\theta \in \Theta_{\text{lex}}^R$.

For a given set R of strict pairwise preferences and a given model θ , checking whether $\theta \in \Theta^R$ can be evaluated in polynomial time in $|R|$ and in $|\theta|$ by solving the following linear program \mathcal{P}_θ , where there is one variable $e_{(A,B)} \geq 0$ for each pair (A, B) in R :

$$(\mathcal{P}_\theta^R) \begin{cases} \min \sum_{(A,B) \in R} e_{(A,B)} \\ \sum_{S \in \theta} (I_A(S) - I_B(S)) v_S + e_{(A,B)} \geq 1 & \forall (A, B) \in R, \\ e_{(A,B)} \geq 0 & \forall (A, B) \in R, \\ v_S \in \mathbb{R} & \forall S \in \theta. \end{cases}$$

We have that $\theta \in \Theta^R$ if and only if the optimal value of \mathcal{P}_θ^R is 0, because in this case we can find values for variables v_S that respect all the preferences in R without the help of the additional slack variables $e_{(A,B)}$.

We now show how to derive, from \mathcal{P}_θ^R , a MIP formulation for computing a model $\theta \in \Theta_{\text{lex}}^R$. For this, we first compute $\text{deg}(R) = \min\{\text{deg}(\theta) : \theta \in \Theta^R\}$, by using Algorithm 1. We then add a binary variable b_S for each $S \in [\mathcal{F}]^{\text{deg}(R)}$, as well as big-M constraints to ensure that $b_S = 1$ iff $S \in \theta$ (i.e., S belongs to the model $\theta \in \Theta_{\text{lex}}^R$). Determining a model $\theta \in \Theta_{\text{lex}}^R$ can be done by solving the following lexicographic optimization problem:

$$(\mathcal{P}_{\text{lex}}^R) \begin{cases} \min \text{lex} \sum_{S \in [\mathcal{F}]^{\text{deg}(R)}} b_S, \sum_{S \in [\mathcal{F}]^{\text{deg}(R)}} b_S |S| \\ \sum_{S \in [\mathcal{F}]^{\text{deg}(R)}} (I_A(S) - I_B(S)) v_S \geq 1 & \forall (A, B) \in R, \\ -b_S M \leq v_S \leq b_S M & \forall S \in [\mathcal{F}]^{\text{deg}(R)}, \\ b_S \in \{0, 1\} & \forall S \in [\mathcal{F}]^{\text{deg}(R)}. \end{cases} \tag{1}$$

where $M = (2^{\sum_{i=1}^{\text{deg}(R)} \binom{n}{i}} + |R|) \times (|R|)^{2|R|+2}$, so that if the values v_S can be set to satisfy constraints 1, then there exist such values in the interval $[-M, M]$ (see Papadimitriou, 1981). Every feasible instantiation of variables v_S, b_S in $\mathcal{P}_{\text{lex}}^R$ corresponds to an element $\theta \in \Theta^R$, namely $\theta = \{S \in [\mathcal{F}]^{\text{deg}(R)} : b_S = 1\}$. Lexicographic optimization amounts to determine, among feasible instantiations of v_S, b_S that minimize the first objective $\sum_{S \in [\mathcal{F}]^{\text{deg}(R)}} b_S$, one that minimizes the second objective $\sum_{S \in [\mathcal{F}]^{\text{deg}(R)}} b_S |S|$. It is well-known that this can be achieved as follows (see, e.g., Section 5.1 in Ehrgott, 2005) using a mixed integer programming solver:

- first, we solve the MIP \mathcal{P}_1 obtained by replacing the lexicographic objective function in $\mathcal{P}_{\text{lex}}^R$ by $\min \sum_{S \in [\mathcal{F}]^{\text{deg}(R)}} b_S$;
- denoting by opt_1 the optimal value of \mathcal{P}_1 , we then solve the MIP \mathcal{P}_2 where the objective function in $\mathcal{P}_{\text{lex}}^R$ is replaced by $\min \sum_{S \in [\mathcal{F}]^{\text{deg}(R)}} b_S |S|$, under the additional constraint $\sum_{S \in [\mathcal{F}]^{\text{deg}(R)}} b_S \leq \text{opt}_1$.

As every feasible instantiation corresponds to a model θ of minimal degree $\text{deg}(\theta)$ (i.e., $\text{deg}(\theta) = \text{deg}(R)$), we thus obtain a model $\theta \in \Theta_{\text{lex}}^R$, from which we deduce $(\text{deg}(\theta), \text{card}(\theta), \text{ws}(\theta))$ for $\theta \in \Theta_{\text{lex}}^R$. In the following, we denote by $(\text{deg}_{\text{lex}}, \text{card}_{\text{lex}}, \text{ws}_{\text{lex}})$ the vector $(\text{deg}(\theta), \text{card}(\theta), \text{ws}(\theta))$ for $\theta \in \Theta_{\text{lex}}^R$.

4.4. Determining whether $A \succ_{\sqsubseteq_{1ex}}^R B$

Determining whether $A \succ_{\sqsubseteq_{1ex}}^R B$ amounts to solve:

$$(\mathcal{P}_{A \succ_{\sqsubseteq_{1ex}}^R B}) \left\{ \begin{array}{l} \min \sum_{S \in [\mathcal{F}]^{\deg(R)}} b_S |S| \\ \sum_{S \in [\mathcal{F}]^{\deg(R)}} b_S \leq \text{card}_{1ex}, \quad (a) \\ \sum_{S \in [\mathcal{F}]^{\deg(R)}} (I_B(S) - I_A(S))v_S \geq 0, \quad (b) \\ \sum_{S \in [\mathcal{F}]^{\deg(R)}} (I_X(S) - I_Y(S))v_S \geq 1 \quad \forall (X, Y) \in R. \quad (c) \\ -b_S M \leq v_S \leq b_S M \quad \forall S \in [\mathcal{F}]^{\deg(R)}, \\ b_S \in \{0, 1\} \quad \forall S \in [\mathcal{F}]^{\deg(R)}. \end{array} \right. \quad (2)$$

A feasible solution of $\mathcal{P}_{A \succ_{\sqsubseteq_{1ex}}^R B}$ yields a model θ satisfying $\deg(\theta) = \deg_{1ex}$ (variables b_S are only defined for $S \in [\mathcal{F}]^{\deg(R)}$) and $\text{card}(\theta) = \text{card}_{1ex}$ (by constraint (2a)) on the value of $\text{card}(\theta)$). Furthermore, constraint (2b) ensures that $\theta \in \Theta_{B \succeq A}^R$, while constraint (2c) ensures that $\theta \in \Theta^R$. If the optimal value of $\mathcal{P}_{A \succ_{\sqsubseteq_{1ex}}^R B}$ is ws_{1ex} , then the corresponding model θ belongs to $\theta \in \Theta_{\sqsubseteq_{1ex}}^R$ (because then $(\deg(\theta), \text{card}(\theta), \text{ws}(\theta)) = (\deg_{1ex}, \text{card}_{1ex}, \text{ws}_{1ex})$), and thus there exists $\theta \in \Theta_{\sqsubseteq_{1ex}}^R \cap \Theta_{B \succeq A}^R$. Consequently:

- if the optimal value is strictly greater than ws_{1ex} , or the polyhedron is empty, then $\Theta_{\sqsubseteq_{1ex}}^R \cap \Theta_{B \succeq A}^R = \emptyset$ and hence $A \not\succeq_{\sqsubseteq_{1ex}}^R B$ (by Proposition 1);
- if the optimal value of $\mathcal{P}_{A \succ_{\sqsubseteq_{1ex}}^R B}$ is ws_{1ex} , then $A \not\succeq_{\sqsubseteq_{1ex}}^R B$.

5. Numerical tests

We call hereafter OROR the learning approach consisting in computing $(\deg(R), \text{card}(R), \text{ws}(R))$ and using $\succ_{\sqsubseteq_{1ex}}^R$ for preference prediction, for Occam’s Robust Ordinal Regression. Numerical tests were carried out on Google Colab,⁶ to compare OROR with state-of-the-art approaches in two different settings:

- A first set of experiments were carried out on *synthetic data*, i.e., obtained by simulating a user. They aimed at evaluating our approach in an ideal setting where a θ -additive model perfectly fits the preferences.
- A second set of experiments were carried out on real-world data for content-based filtering methods (more precisely, movies described by binary attributes). These tests aimed at evaluating how our approach deals with partially described alternatives (i.e., with possible “collisions” if two distinct alternatives share the same description), compared to other state of the art approaches.

In both sets of experiments, we start with a learning set of preferences, consisting of a strict weak order on a subset of alternatives. Based on this learning set, pairwise preference predictions are then requested on random pairs of alternatives (pairs not in the learning set). As said earlier, the model may reject a prediction if it is not robust enough given the available preference data (i.e., if there is no \sqsubseteq_{1ex} -ordinal dominance).

5.1. The synthetic and real-world datasets

The dataset consists of ratings assigned by a user (DM) on a set \mathcal{A} of N alternatives. Given a set $\mathcal{F} = \{a_1, \dots, a_n\}$ of binary features, a learning set \mathcal{A}_{train} consists of $k \leq N$ ratings of alternatives in \mathcal{A} , where each alternative A_i ($i = 1, \dots, N$) is described by a binary vector $\bar{A}_i = (A_i^1, \dots, A_i^n)$, with $A_i^j = 1$ if $a_j \in A_i$, and $A_i^j = 0$ otherwise. The user rating of A_i is denoted by r_i . The set of known strict preferences is $R = \{(A_i, A_j) \in \mathcal{A}_{train}^2 : r_i > r_j\}$. We recall that alternatives A_i and A_j

are *incomparable* (and not indifferent) if $r_i = r_j$. The ratings thus induce a *strict weak order* on \mathcal{A}_{train} .

The *real-world data* consist of ratings of movies by users picked up from the IMDb dataset.⁷ This is a dataset of movie reviews that contains over 50k reviews. Each movie A_i is described by a set of binary features A_i^j , and the ratings r_i are integer values ranging from 1 to 10. The experiments were conducted with a dataset of 50 users (randomly sampled) who each rated at least $k = 100$ movies. Each movie is described using a subset of $n = 8$ binary features (corresponding to the main genres of the movie, e.g., “adventure”, “animation”, “children”, “comedy”, “fantasy”, etc.). In the dataset, 18% (resp. 29%, 32%, 13%, 4%, 1%) of the movies have 1 (resp. 2, 3, 4, 5, 6) positive binary feature(s). This implies that $\deg(R) \leq 6$.

The *synthetic data* consists of a randomly sampled θ -additive function $f_{\theta,v}$ where the value function v is integer-valued. The rating of an alternative A_i corresponds to the value $f_{\theta,v}(A_i)$. For sampling a function $f_{\theta,v}$, we first sample a set θ and then we sample parameters v_S for $S \in \theta$. More precisely, the generation of θ is achieved as follows. First, θ is initialized as the set of singletons $\{a_1\}, \{a_2\}, \dots, \{a_n\}$, then we add $\lfloor \alpha \times (2^n - n) \rfloor$ 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, which 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 sets S we sample is $\mathbb{E}[|S|] = 2 + (1-p)(1-p)^{n-1}/p$. Once θ is set, we sample the parameters v_S for each $S \in \theta$ with a discrete uniform distribution $\mathcal{U}\{-2, 2\}$ (where all integers in $\{-2, \dots, 2\}$ are equally probable). The sampling of $f_{\theta,v}$ thus depends on two parameters p and α . In the tests, p varies in $[0.1, 0.9]$ and α in $[0.1, 0.5]$.

5.2. Baseline models

We briefly describe here the baseline models to which OROR is compared. Throughout the subsection, we have $\theta = [\mathcal{F}]^{\deg(R)}$ and each alternative A is described by an *augmented* binary vector $\bar{A}_\theta = (I_A(S_1), \dots, I_A(S_{|\theta|}))$, where $S_1, \dots, S_{|\theta|}$ are the subsets of \mathcal{F} of size less than or equal to $\deg(R)$. The three baseline models (OR, SVM and GP) we consider involve a reject option (either inherently or because we introduce it). We describe them below.

Robust ordinal regression (ROR). This corresponds to the robust ordinal regression model proposed by Greco et al. (2008). It consists in predicting a preference between two alternatives A, B if $f_{\theta,v}(A) > f_{\theta,v}(B)$ for each $v \in V_\theta^R$. Note that $\theta = [\mathcal{F}]^{\deg(R)}$ implies that $V_\theta^R \neq \emptyset$. We use Algorithm 1 to compute $\deg(R)$.

Support vector machine (SVM). This baseline model is inspired by the 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.⁸ In our setting, each couple $(A, B) \in R$ yields an example formalized by a $|\theta|$ -dimensional vector $\bar{A}_\theta - \bar{B}_\theta$ labeled by 1 as $A > B$. The model fitted using the deterministic SVM is a hyperplane $w \in \mathbb{R}^{|\theta|}$ that separates all

⁷ www.kaggle.com/datasets/gauravduttakiit/imdb-recommendation-engine.

⁸ We use the SVC implementation from the scikit-learn python library (Pedregosa et al., 2011).

⁶ two virtual CPU at 2.2 GHz, 13 GB RAM.

vectors $(\bar{A}_\theta - \bar{B}_\theta)$ from vectors $(\bar{B}_\theta - \bar{A}_\theta)$. This hyperplane is obtained by solving the following optimization problem:

$$\min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^T \mathbf{w} \tag{3}$$

subject to:

$$\mathbf{w} \cdot (\bar{A}_\theta - \bar{B}_\theta) \geq 1, \forall (A, B) \in R \tag{4}$$

Since the alternatives are projected using $\theta = [\mathcal{F}]^{\text{deg}(R)}$, the problem is linearly separable; that is, there always exists a hyperplane that satisfies the constraints of the optimization problem. To predict a preference on a pair A, B of alternatives, we determine the side of the hyperplane on which the vector $\bar{A}_\theta - \bar{B}_\theta$ lies by evaluating the SVM score $\mathbf{w} \cdot (\bar{A}_\theta - \bar{B}_\theta)$. If this value is strictly positive (resp. strictly negative), then we predict $A > B$ (resp. $B > A$). To introduce a reject option, Platt scaling (Platt, 1998) is used to transform the SVM scores into probabilities $\mathbb{P}(A > B|R)$. This probability is obtained through a sigmoid function parameterized by parameters $\alpha, \beta \in \mathbb{R}^2$:

$$\mathbb{P}_{\alpha, \beta}(A > B|R) = \frac{1}{1 + \exp(-\alpha[\mathbf{w} \cdot (\bar{A}_\theta - \bar{B}_\theta)] - \beta)}$$

These parameters are fitted using maximum likelihood estimation:

$$\hat{\alpha}, \hat{\beta} = \arg \max_{\alpha, \beta \in \mathbb{R}} \sum_{(A, B) \in R} \log(\mathbb{P}_{\alpha, \beta}(A > B|R))$$

The optimal parameters $\hat{\alpha}$ and $\hat{\beta}$ are then used to compute the probability $\mathbb{P}(A > B|R) = \mathbb{P}_{\hat{\alpha}, \hat{\beta}}(A > B|R)$. We predict $A > B$ if $\mathbb{P}(A > B|R) \geq t$ for a given threshold t . The higher the threshold, the more rejections are made.

Gaussian process (GP). We consider the preference learning method proposed by Chu and Ghahramani (2005). Given a set R of known pairwise preferences, the GP approach to preference learning approximates a probability distribution $\mathbb{P}(f|R)$ on the space f of utility functions defined on \mathcal{A} .⁹ Formulated in our setting, the idea is to assume that each utility value $f(A)$ is a random variable following a normal distribution. The joint distribution on \mathcal{A} , the so-called Gaussian process, is a multivariate Gaussian, characterized by a mean function and a covariance matrix. The approach is *non-parametric*, in the sense that no analytical definition of f is used, but only the idea that alternatives A, B with similar vectors \bar{A}_θ and \bar{B}_θ should have similar utilities $f(A)$ and $f(B)$. This is formalized by defining the covariance matrix using a kernel function (Schölkopf & Smola, 2002), i.e., a symmetric function $k : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}$ such that $k(A, B) \geq 0$ reflects the similarity between \bar{A}_θ and \bar{B}_θ . More precisely, we use a Mercer kernel function $k(A, B) = \exp\left(-\frac{1}{2}(\bar{A}_\theta - \bar{B}_\theta)^T(\bar{A}_\theta - \bar{B}_\theta)\right)$. Given a learning set $\mathcal{A}_{\text{train}} = \{A_1, \dots, A_k\}$, the prior probability of the utility values $f(A_1), \dots, f(A_k)$ is a zero mean multivariate Gaussian:

$$\mathbb{P}(f) = \frac{1}{(2\pi)^{\frac{k}{2}} |\Sigma_0|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} f^T \Sigma_0^{-1} f\right)$$

where $f = [f(A_1), f(A_2), \dots, f(A_k)]^T$ and:

$$\Sigma_0 = \begin{bmatrix} k(A_1, A_1) & \dots & k(A_1, A_k) \\ \vdots & \ddots & \vdots \\ k(A_k, A_1) & \dots & k(A_k, A_k) \end{bmatrix}.$$

Note that the terms on the diagonal of the matrix are all equal to 1. A probit likelihood function is used to update the prior from the set R of pairwise preferences inferred from $\mathcal{A}_{\text{train}}$. The posterior distribution $\mathbb{P}(f|R)$ is approximated, using Laplace’s approximation, by a new multivariate Gaussian distribution centered on f_{MAP} , the maximum a posterior estimate of f (obtained by solving a convex programming problem) (see Brochu, Cora, & De Freitas, 2010, Section 3.1).

⁹ We use the GaussianProcessClassifier implementation from the scikit-learn python library (Pedregosa et al., 2011).

Predicting a preference on a pair A, B of alternatives works as follows. The following prior joint multivariate Gaussian distribution is considered:

$$\begin{bmatrix} f \\ f_{AB} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \Sigma_0 & K_{AB} \\ K_{AB}^T & \Sigma_{AB} \end{bmatrix}\right)$$

where $f_{AB} = [f(A), f(B)]^T$ and:

$$K_{AB} = \begin{bmatrix} k(A_1, A) & \dots & k(A_k, A) \\ k(A_1, B) & \dots & k(A_k, B) \end{bmatrix}^T \text{ and } \Sigma_{AB} = \begin{bmatrix} k(A, A) & k(A, B) \\ k(B, A) & k(B, B) \end{bmatrix}.$$

The predictive preference $\mathbb{P}(A > B|R)$ is evaluated from the predictive distribution $\mathbb{P}(f_{AB}|R)$, which is itself computed from the (approximated) posterior $\mathbb{P}(f|R)$. It can be shown (see Chu & Ghahramani, 2005) that $\mathbb{P}(A > B|R)$ corresponds to:

$$\mathbb{P}(A > B|R) = \Phi\left(\frac{\mu_A^* - \mu_B^*}{\sigma^*}\right)$$

where Φ is the cumulative distribution function of the standard normal distribution, $[\mu_A^*, \mu_B^*]^T = K_{AB}^T \Sigma_0^{-1} f_{\text{MAP}}$ and σ^* is computed from $\Sigma_0, \Sigma_{AB}, K_{AB}$ and the variance of the probit model (for the detailed formula, see Chu & Ghahramani, 2005, Section 2.3). As for SVM, we predict $A > B$ if $\mathbb{P}(A > B|R) \geq t$ for a given threshold t . The higher the threshold, the more rejections are made.

5.3. Experimental setup

In all experiments, the dataset is a set \mathcal{A} of N alternatives, described by a set \mathcal{F} of n binary features, and an associated rating vector r (integer values). The rating $r(A)$ of each alternative $A \in \mathcal{A}$ is known. To compare the performances of the different learning methods, we extract a subset $\mathcal{A}_{\text{train}}$ of k alternatives from \mathcal{A} , on which the models are trained. The alternatives in $\mathcal{A}_{\text{train}}$ are chosen uniformly at random. We then randomly sample 100 pairs $\{A, B\}$ in \mathcal{A} such that $A \notin \mathcal{A}_{\text{train}}$ or $B \notin \mathcal{A}_{\text{train}}$ (possibly neither A nor B belongs to $\mathcal{A}_{\text{train}}$), and we compare the predicted pairwise preference with the actual preference: $A > B$ if $r(A) > r(B)$, $B > A$ if $r(B) > r(A)$, $A \sim B$ (incomparability) if $r(A) = r(B)$. The extraction of a subset $\mathcal{A}_{\text{train}}$ from \mathcal{A} , the training of each model and the (100) pairwise preference predictions are performed 10 times, and the prediction performances are averaged over the 10 runs. We detail below the parameters that are used for the experiments on synthetic data and for the experiments on real-world data.

Synthetic data. The experiments on synthetic data were conducted with $|\mathcal{F}| = 8$ binary features, which yields a set \mathcal{A} of $2^{|\mathcal{F}|} = 256$ alternatives, and the set of parameters $(\alpha, p) = (0.1, 0.9)$ for the generation of $f_{\theta, v}$. This set of parameters yields functions $f_{\theta, v}$ that are usually up to 4-additive, with an average $|\theta|$ equal to 12. This setting is not really restrictive as, given the number of strict pairwise preferences in R that are considered in our experiments (i.e., $|R| \leq \binom{|\mathcal{A}_{\text{train}}|}{2}$), it is unlikely that R cannot be represented by using a function $f_{\theta, v}$ of degree up to 4. The size of $\mathcal{A}_{\text{train}}$ indeed varies between 4 and 18, from which between $|R| = \binom{4}{2} = 6$ and $\binom{18}{2} = 153$ pairwise preferences can be inferred.

Real-world data. For each of the 50 users that have rated at least 100 movies, a dataset \mathcal{A} including between 45 and 100 alternatives is first extracted. A training set $\mathcal{A}_{\text{train}}$ is then extracted from \mathcal{A} , with $|\mathcal{A}_{\text{train}}|$ corresponding to 90% of $|\mathcal{A}|$ (which is common practice in machine learning, in particular for performing 10-fold cross-validation). The size of $\mathcal{A}_{\text{train}}$ thus varies from 5 to 10, from which between $|R| = \binom{5}{2} = 10$ and $\binom{10}{2} = 45$ pairwise preferences can be inferred.

5.4. Evaluation metrics

We outline here the specific metrics that will be used to evaluate the OROR approach and compare it to other methods. To define our metrics we consider the 9 cases that can occur in the confusion matrix defined below.

Table 2

Confusion matrix.

Predicted/Real	(B)etter ($r(A) > r(B)$)	(W)orse ($r(B) > r(A)$)	(I)ncomparable ($r(A) = r(B)$)
(B)etter ($A > B$)	BB	BW	BI
(W)orse ($B > A$)	WB	WW	WI
(R)ejection	RB	RW	RI

Confusion matrix. For a given pair of alternatives $(A, B) \in \mathcal{A}^2$ each model could either infer (predicted output) that A is *better* than B ($A > B$), or that A is *worse* than B ($B > A$), or it could *reject* the instance (no prediction). Then, as outlined earlier, by comparing $r(A)$ and $r(B)$, we can have (real outputs) that A is indeed better than B if $r(A) > r(B)$ or that A is worse than B if $r(B) > r(A)$ or that A and B are incomparable if $r(A) = r(B)$. Our metrics are based on the *confusion matrix* defined in Table 2, where the rows symbolizes the predicted outputs and the columns the real outputs.

Precision, recall and F-score are widely used metrics to assess a (binary) classification method with a positive class and a negative class. In the standard setting, they are only four components in the confusion matrix: True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN). For instance, a true (resp. false) positive corresponds to an instance that is predicted as positive and that is indeed positive (resp. is negative). Before detailing how we adapt the metrics to our setting where the confusion matrix has nine components (BB, BW, etc.), we recall how the metrics are defined from TP, TN, FP and FN in the standard setting (see e.g., Goutte & Gaussier, 2005; Van Rijsbergen, 1979):

Precision. The precision is defined as the ratio between the number of true positives among all the instances predicted as positive.

$$P = \frac{TP}{TP + FP}.$$

Recall. The recall is defined as the ratio between the number of true positives among all the positive instances.

$$R = \frac{TP}{TP + FN}.$$

Both metrics are complementary because the precision metric penalizes the models making unreliable positive predictions, while the recall metric penalizes the models that avoid making positive predictions.

These classic metrics have been adapted to various contexts, where TP, TN, FP and FN need to be redefined. This is, for instance, the case when the classifier may predict a subset of classes (Del Coz, Díez, & Bahamonde, 2009; Zaffalon, Corani, & Mauá, 2012). We adapt the metrics in a similar, yet different, manner: semantically, rejecting a pair A, B , i.e., not predicting $A > B$ nor $B > A$, is indeed different from predicting both $A > B$ and $B > A$. Other papers have defined metrics to evaluate classifiers using reject options (Hendrickx et al., 2024). However, these metrics require knowing which class would have been predicted, if the sample were not rejected. Such information is not available in our method, which leads us to define our own metrics.

To adapt these metrics to our confusion matrix, we now explicit how we define TP, FP, TN and FN from the components of the matrix in Table 2.¹⁰ In our setting, the positive instances correspond to pairs A, B such that $r(A) \neq r(B)$ (i.e., there is a strict preference between A and B), and the negative instances to pairs A, B such that $r(A) = r(B)$ (i.e., A and B are incomparable). A true positive corresponds then to making a prediction when $r(A) \neq r(B)$. As these predictions may be correct or wrong, we denote by CTP (resp. WTP) the number of correct (resp.

¹⁰ Another approach could be to consider W , B , and I as three different classes and to resort to metrics for multi-class classification (Grandini, Bagli, & Visani, 2020). However, this approach would not capture the symmetry between cases WW and BB , nor between WB and BW .

wrong) true positives. A correct true positive corresponds to predicting $A > B$ (resp. $B > A$) when $r(A) > r(B)$ (resp. $r(B) > r(A)$), while a wrong true positive corresponds to predicting $A > B$ (resp. $B > A$) when $r(B) > r(A)$ (resp. $r(A) > r(B)$). Conversely, a false positive corresponds to predicting $A > B$ or $B > A$ while $r(A) = r(B)$. A true negative corresponds to rejecting a pair A, B when $r(A) = r(B)$. Conversely, a false negative corresponds to rejecting a pair A, B while $r(A) \neq r(B)$. This yields the following equations:

$$CTP = BB + WW, \quad WTP = BW + WB, \quad TN = RI,$$

$$FP = BI + WI, \quad FN = RW + RB.$$

Precision and recall are thus defined as follows in our setting:

$$P = \frac{CTP}{CTP + WTP + FP}, \quad R = \frac{CTP}{CTP + WTP + FN}.$$

The precision indicates the proportion of correct comparisons among all the preference predictions made by the model, while the recall indicates the proportion of correct comparisons among all the comparable pairs.

F_β -Score. F_β -score is a metric that combines precision and recall to provide a balanced evaluation of a model’s performance. It is obtained by computing the weighted harmonic mean of precision and recall:

$$F_\beta = (1 + \beta^2) \times \frac{P \times R}{(\beta^2 \times P) + R}.$$

The F_β -score, characterized by its tunable β parameter, offers a flexible approach to balancing precision and recall. A β value of 1 indicates an equal emphasis on both precision and recall, maintaining a harmonious balance. When β exceeds 1, the score leans towards prioritizing recall, emphasizing the identification of all positive instances. Conversely, a β value less than 1 shifts the focus towards precision, stressing the importance of correctly identifying positive instances while minimizing false positives.

5.4.1. Results on synthetic data

In Table 3, we give a comprehensive assessment of OROR and the baseline models (using thresholds $t \in \{0.5, 0.6, 0.8, 0.9\}$ for SVM and GP) across 300 couples $(f_{\theta,v}, \mathcal{A}_{train})$, where $f_{\theta,v}$ is a randomly sampled integer-valued θ -additive function (as described in Section 5.1) and \mathcal{A}_{train} is a subset of $k \in \{4, \dots, 18\}$ alternatives (random subsets of \mathcal{F}). The functions $f_{\theta,v}$ are sampled with parameters $(n, p, \alpha) = (6, 0.2, 0.2)$. The evaluation metrics include precision, recall, and F_1 -score (see Section 5.4). From each couple $(f_{\theta,v}, \mathcal{A}_{train})$, we define a set R of strict pairwise preferences, by comparing each pair of alternatives in \mathcal{A}_{train} . The size $|R|$ ranges from 6 to 153. The preference sets R are categorized into four groups according to $|R|$: [6, 40], [40, 80], [80, 120], and [120, 153].

In Fig. 2, the average F_β -scores for each $|R|$ group, and a confidence interval of 95% around them, are given for OROR and for the baseline models, for various β values. For the ROR and SVM models, a threshold $t = 0.6$ is used in this figure because preliminary tests showed that this value of t achieves the best trade-off between precision and recall.

The information presented in Table 3 helps understand the impact of robustness requirements on model performances. Obviously, for the SVM and GP models, the recall increases and the precision decreases as the threshold increases. The best compromise according to the F_1 -score is to favor smaller thresholds for $|R| \in [6, 80]$ and larger thresholds for $|R| \in [80, 153]$. Additionally, while ROR is always the most accurate model, its low recall undermines its precision advantage, resulting in lower F_1 -scores for $|R| \in [6, 80]$. The OROR model strikes a good balance between prediction quantity and quality for $|R| \in [80, 153]$, outperforming the other models in terms of F_1 -score for $|R| \geq 80$, and significantly surpassing them for $|R| \geq 120$. This advantage for $|R| \geq 80$ is further demonstrated in Fig. 2, where F_β -scores for $\beta \in \{0, \dots, 4\}$, quantifying various trade-offs between accuracy and recall, are compared across the different models. The OROR model often outperforms the baseline models, especially the ROR model for $\beta \leq 1.5$.

Table 3

Performance comparison between OROR and the baseline models on synthetic data for $n = 6$, $p = 0.2$ and $\alpha = 0.2$.

Model	t	SVM				GP				ROR	OROR
		0.5	0.6	0.8	0.9	0.5	0.6	0.8	0.9		
$R \in [6, 40]$	P	0.57	0.62	0.64	0.72	0.59	0.63	0.72	0.81	0.84	0.67
	R	0.69	0.57	0.51	0.31	0.63	0.56	0.29	0.13	0.16	0.47
	F_1	0.62	0.59	0.57	0.46	0.61	0.59	0.45	0.33	0.37	0.55
$R \in [40, 80]$	P	0.71	0.79	0.81	0.85	0.70	0.76	0.88	0.94	0.95	0.83
	R	0.87	0.80	0.75	0.63	0.84	0.77	0.51	0.32	0.39	0.74
	F_1	0.78	0.78	0.76	0.71	0.77	0.76	0.63	0.50	0.55	0.76
$R \in [80, 120]$	P	0.74	0.82	0.84	0.87	0.74	0.82	0.93	0.97	1.00	0.85
	R	0.92	0.85	0.82	0.75	0.90	0.89	0.62	0.44	0.57	0.90
	F_1	0.81	0.82	0.82	0.78	0.81	0.85	0.73	0.60	0.71	0.86
$R \in [120, 153]$	P	0.80	0.88	0.90	0.88	0.83	0.87	0.92	1.00	1.00	0.96
	R	0.96	0.93	0.87	0.83	0.96	0.92	0.69	0.50	0.79	0.98
	F_1	0.86	0.89	0.87	0.85	0.89	0.89	0.76	0.64	0.87	0.96




Fig. 2. Evolution of the F_β score w.r.t. β on synthetic data.

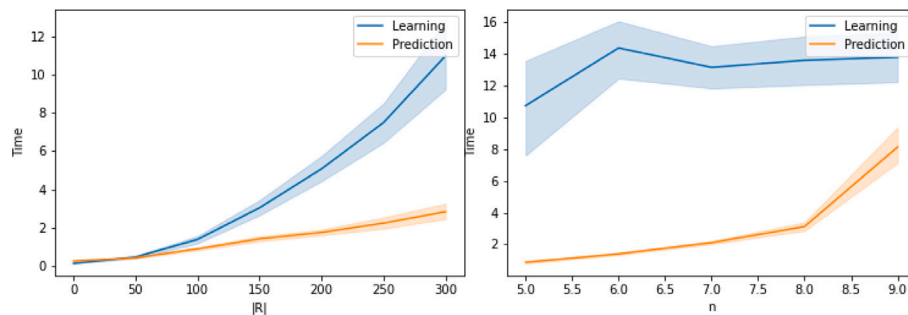


Fig. 3. Running times of OROR (in seconds).

The curves in Fig. 3 give the average running times of OROR (in seconds, averaged over 20 instances) according to the number n of features (for $300 \leq |R| \leq 400$) and the number $|R|$ of known strict pairwise preferences (for $n = 8$). The orange curve gives the average running time for *one* pairwise preference prediction; this is the most time-consuming phase: note indeed that learning ($\text{deg}(R), \text{card}(R), \text{ws}(R)$) is only performed once for each R . In contrast, 100 preference predictions are made for each R in our tests.

5.4.2. Results on real-world data

Fig. 4 gives the average F_β -score (for $\beta \in [0, 4]$) and a 95% confidence interval for both OROR and the baseline models across 50 random users for different sizes $|R|$, evaluated through 10-fold validation (Ojala & Garriga, 2010).

The first observation that can be made is that, for $|R| \in [60, 160]$, the best F_1 -score is achieved by ROR, after using Algorithm 1 to compute the degree $\text{deg}(R)$ of the preference set R . Nevertheless, for $|R| \in [5, 60]$ (i.e., when there are fewer preference examples), the OROR model is the only one to outperform SVM and GP for all β values. More generally, the results seem to confirm that the OROR model outperforms overall SVM and GP. This makes the OROR model particularly

valuable for addressing the cold-start challenges often encountered in recommendation systems.

6. Conclusion

We have presented here a robust ordinal regression method for subsets comparisons with interactions. The model we use is not restrictive, in the sense that any strict weak order on subsets can be represented. The learning method achieves a trade-off between the number of predicted preferences and the accuracy of the predictions, by relying on a new ordinal dominance relation between subsets.

Several research directions are worth investigating, notably the trade-off that should be made between the flexibility of the decision model and the ability to detect mistakes. The model we propose here is indeed very flexible and may allow any transitive preferences, which makes it compatible with preference examples that might almost seem irrational. While there may be good reasons for the DM to express such preferences, it is also well-known that mistakes are often made when providing preference statements. In such a case, the model might overfit to accommodate these mistakes. A way to overcome this drawback is to use a probabilistic approach for eliciting the parameter set, e.g., by Bayesian updating (Gilbert, Ouaguenouni, Öztürk, & Spanjaard, 2023).




Fig. 4. Evolution of the F_β score w.r.t. β on real-world data.

The realized trade-off (i.e., a trade-off favoring flexibility or robustness to errors) is tuned by the choice of the prior on the parameter set. Another research direction is the adaptation of the approach to an active learning setting where one interactively determines a sequence of queries to minimize the cognitive burden for the decision maker or a better consideration of potential “errors” in the preferences used as a learning set.

CRedit authorship contribution statement

Hugo Gilbert: Writing – original draft, Writing – review & editing. **Mohamed Ouaguenouni:** Writing – original draft, Writing – review & editing. **Meltem Öztürk:** Writing – original draft, Writing – review & editing. **Olivier Spanjaard:** Writing – original draft, Writing – review & editing.

Acknowledgments

We acknowledge the support of the French Agence Nationale de la Recherche (ANR), under grant ANR20-CE23-0018 (project THEMIS). We would like to thank the reviewers for their thoughtful comments that helped us substantially improve the manuscript, as well as Yann Chevaleyre for insightful discussions on the efficient computation of the degree of a set of pairwise preferences.

Appendix. Properties of the θ -ordinal dominance relation

Proposition 2. *The following properties hold for $>_\theta^R$:*

- (i) $>_\theta^R$ is asymmetric.
- (ii) $>_\theta^R$ may not be complete.
- (iii) $>_\theta^R$ is not necessarily negatively-transitive.

Proof. (i) $A >_\theta^R B \Rightarrow \forall v \in V_\theta^R, f_{\theta,v}(A) > f_{\theta,v}(B)$. Thus there is no value function $v' \in V_\theta^R$ such that $f_{\theta,v'}(A) < f_{\theta,v'}(B)$.

(ii) As shown in Example 2, we may have $v, v' \in V_\theta^R$ such that $f_{\theta,v}(A) > f_{\theta,v}(B)$ and $f_{\theta,v'}(B) > f_{\theta,v'}(A)$. We have then neither $A >_\theta^R B$ nor $B >_\theta^R A$, and thus $>_\theta^R$ may not be complete.

(iii) Let $\mathcal{F} = \{a_1, a_2, a_3\}$, $R = \{(\{a_1\}, \{a_3\})\}$, $\theta = \{\{a_1\}, \{a_2\}, \{a_3\}\}$ and v, v' two value functions defined as follows:

$$v(\{a_1\}) = 2, v(\{a_2\}) = 3, v(\{a_3\}) = 1,$$

$$v'(\{a_1\}) = 3, v'(\{a_2\}) = 1, v'(\{a_3\}) = 2.$$

We have that $v, v' \in V_\theta^R$ as $f_{\theta,v}(\{a_1\}) > f_{\theta,v}(\{a_3\})$ and $f_{\theta,v'}(\{a_1\}) > f_{\theta,v'}(\{a_3\})$.

It follows from $f_{\theta,v}(\{a_2\}) > f_{\theta,v}(\{a_1\})$ that $\neg(\{a_1\} >_\theta^R \{a_2\})$.

It follows from $f_{\theta,v'}(\{a_3\}) > f_{\theta,v'}(\{a_2\})$ that $\neg(\{a_2\} >_\theta^R \{a_3\})$.

Yet $\{a_1\} >_\theta^R \{a_3\}$ by definition of R . \square

Proposition 3. *Given a set R of strict pairwise comparisons, and $\theta \in \Theta^R$, if $R' \subseteq R$ then: (i) $\theta \in \Theta^{R'}$; (ii) $A >_\theta^R B \Rightarrow A >_{\theta'}^R B$; (iii) $A >_{\theta'}^R B \Rightarrow \neg(B >_\theta^R A)$.*

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 conclusion as $\Theta^R \subseteq \Theta^{R'}$ (because R contains all the preference constraints in R' , along with additional constraints).

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

Proposition 4. *Let $\theta, \theta' \in \Theta^R$. If $\theta' \subseteq \theta$, then the following assertions hold:*

- (i) $A >_\theta^R B \Rightarrow A >_{\theta'}^R B$; (ii) $A \sim_\theta^R B \Rightarrow A \sim_{\theta'}^R B$; (iii) $A >_{\theta'}^R B \Rightarrow \neg(B >_\theta^R A)$.

Proof. (i) is true because if $f_{\theta,v}(A) > f_{\theta,v}(B)$ for all $v \in V_\theta^R$, then we should also have $f_{\theta',v}(A) > f_{\theta',v}(B)$ for all $v \in V_{\theta'}^R$. Indeed, each element of $V_{\theta'}^R$ can be seen as a value function in V_θ^R in which the parameters v_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: $B >_{\theta'}^R A \Rightarrow B >_\theta^R A$ by (i), and $B >_\theta^R A \Rightarrow \neg(A >_\theta^R B)$ because strict preferences are asymmetric. \square

References

Adam, L., & Destercke, S. (2021). Possibilistic preference elicitation by minimax regret. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021* (pp. 718–727).

Andreopoulou, Z., Koliouka, C., & Zopounidis, C. (2017). Multicriteria and Clustering: Classification Techniques in Agrifood and Environment. In *Cooperative Management*. Springer.

Angilella, S., Corrente, S., & Greco, S. (2015). Stochastic Multiobjective Acceptability Analysis for the Choquet integral preference model and the scale construction problem. *European Journal of Operational Research*, 240(1), 172–182.

Angilella, S., Corrente, S., Greco, S., & Slowinski, R. (2016). Robust Ordinal Regression and Stochastic Multiobjective Acceptability Analysis in Multiple Criteria Hierarchy Process for the Choquet integral preference model. *Omega*, 63, 154–169.

Bartee, E. M. (1971). Problem solving with ordinal measurement. *Management Science*, 17(10), B–622.

Benabbou, N., Leroy, C., Lust, T., & Perny, P. (2021). Combining preference elicitation with local search and greedy search for matroid optimization. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021* (pp. 12233–12244). AAAI Press.

Bigot, D., Fargier, H., Mengin, J., & Zanuttini, B. (2012). Using and learning GAI-decompositions for representing ordinal rankings. In *ECAI'2012 Workshop on Preference Learning (PL 2012)* (pp. 5–10).

Blumer, A., Ehrenfeucht, A., Haussler, D., & Warmuth, M. K. (1987). Occam's razor. *Information Processing Letters*, 24(6), 377–380.

Bourdache, N., Perny, P., & Spanjaard, O. (2019). Incremental elicitation of rank-dependent aggregation functions based on Bayesian linear regression. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019* (pp. 2023–2029).

Boutillier, C., Patrascu, R., Poupart, P., & Schuurmans, D. (2006). Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 170(8), 686–713.

Bresson, R., Cohen, J., Hüllermeier, E., Labreuche, C., & Sebag, M. (2020). Learning 2-additive Hierarchical Choquet Integrals with non-monotonic utilities. In *DA2PL 2020*.

Brochu, E., Cora, V. M., & De Freitas, N. (2010). A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. arXiv preprint arXiv:1012.2599.

- Charnetski, J. R., & Soland, R. M. (1978). Multiple-attribute decision making with partial information: the comparative hypervolume criterion. *Naval Research Logistics Quarterly*, 25(2), 279–288.
- Chu, W., & Ghahramani, Z. (2005). Preference learning with Gaussian processes. In *Proceedings of the Twenty-Second International Conference on Machine Learning, ICML 2005* (pp. 137–144).
- Ciomek, K., Kadziński, M., & Tervonen, T. (2017). Heuristics for prioritizing pair-wise elicitation questions with additive multi-attribute value models. *Omega*, 71, 27–45.
- Corrente, S., Greco, S., Kadziński, M., & Słowiński, R. (2013). Robust ordinal regression in preference learning and ranking. *Machine Learning*, 93(2), 381–422.
- Corrente, S., Greco, S., Kadziński, M., & Słowiński, R. (2014). Robust ordinal regression. *Wiley Encyclopedia of Operations Research and Management Science*, 1–10.
- Corrente, S., & Tasiou, M. (2023). A robust TOPSIS method for decision making problems with hierarchical and non-monotonic criteria. *Expert Systems with Applications*, 214, Article 119045.
- Del Coz, J. J., Díez, J., & Bahamonde, A. (2009). Learning nondeterministic classifiers. *Journal of Machine Learning Research*, 10(10).
- Dias, L. C., Morton, A., & Quigley, J. (Eds.). (2018). *Elicitation: The science and art of structuring judgement*. Springer.
- Domshlak, C., & Joachims, T. (2005). Unstructuring user preferences: efficient non-parametric utility revelation. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence, UAI 2005* (pp. 169–177).
- Dragone, P., Teso, S., & Passerini, A. (2017). Constructive preference elicitation over hybrid combinatorial spaces. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, AAAI 2018* (pp. 2943–2950).
- Ehrgott, M. (2005). *Multicriteria optimization*. Springer.
- Fishburn, P. C. (1970). *Utility theory for decision making*. Wiley.
- Fishburn, P. C., & Lavalley, I. H. (1996). Binary interactions and subset choice. *European Journal of Operational Research*, 92, 182–192.
- Fürnkranz, J., & Hüllermeier, E. (2003). Pairwise preference learning and ranking. In *Proceedings of the Fourteenth European Conference on Machine Learning, ECML 2003* (pp. 145–156). Springer.
- Ghaderi, M., Ruiz, F., & Agell, N. (2017). A linear programming approach for learning non-monotonic additive value functions in multiple criteria decision aiding. *European Journal of Operational Research*, 259(3), 1073–1084.
- Ghosh, S., & Kalagnanam, J. (2003). Polyhedral sampling for multiattribute preference elicitation. In *Proceedings of the Fourth ACM conference on Electronic Commerce, EC 2003* (pp. 256–257).
- Gilbert, H., Ouaguenouni, M., Öztürk, M., & Spanjaard, O. (2023). A hybrid approach to preference learning with interaction terms. In *Proceedings of the Twenty-Sixth European Conference on Artificial Intelligence ECAI 2023* (pp. 835–842). IOS Press.
- Gonzales, C., & Perny, P. (2005). GAI networks for decision making under certainty. In *Proceedings of the First Multidisciplinary Workshop on Advances in Preference Handling, M-PREF 2005* (pp. 100–105).
- Goutte, C., & Gaussier, E. (2005). A probabilistic interpretation of precision, recall and F-score, with implication for evaluation. In *Advances in information retrieval* (pp. 345–359). Springer.
- Grabisch, M. (1996). The application of fuzzy integrals in multicriteria decision making. *European Journal of Operational Research*, 89(3), 445–456.
- Grabisch, M., Kojadinovic, I., & Meyer, P. (2008). 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.
- Grabisch, M., Marichal, J.-L., & Roubens, M. (2000). Equivalent representations of set functions. *Mathematics of Operations Research*, 25(2), 157–178.
- Grandini, M., Bagli, E., & Visani, G. (2020). Metrics for multi-class classification: an overview. arXiv preprint arXiv:2008.05756.
- Greco, S., Mousseau, V., & Słowiński, R. (2008). Ordinal regression revisited: multiple criteria ranking using a set of additive value functions. *European Journal of Operational Research*, 191(2), 416–436.
- Guo, S., & Sanner, S. (2010). Multiattribute Bayesian preference elicitation with pairwise comparison queries. In *Proceedings of the Seventh International Symposium on Neural Networks, ISNN 2010* (pp. 396–403). Springer.
- Hendrickx, K., Perini, L., Van der Plas, D., Meert, W., & Davis, J. (2024). Machine learning with a reject option: A survey. *Machine Learning*, 113(5), 3073–3110.
- Herin, M., Perny, P., & Sokolovska, N. (2023). Learning preference models with sparse interactions of criteria. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023* (pp. 3786–3794).
- Iyengar, V. S., Lee, J., & Campbell, M. (2001). Evaluating multiple attribute items using queries. In *Proceedings of the Third ACM Conference on Electronic Commerce, EC 2001* (pp. 144–153).
- Jacquet-Lagrèze, E., & Siskos, J. (1982). Assessing a set of additive utility functions for multicriteria decision-making, the UTA method. *European Journal of Operational Research*, 10(2), 151–164.
- Kadziński, M. (2022). Robust Ordinal Regression for Multiple Criteria Decision Aiding. In *Intelligent decision support systems: combining operations research and artificial intelligence – Essays in honor of Roman Słowiński* (pp. 185–205). Springer.
- Kompa, B., Snoek, J., & Beam, A. L. (2021). Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digital Medicine*, 4(1).
- Kozlov, M. K., Tarasov, S. P., & Khachiyan, L. G. (1979). Polynomial solvability of convex quadratic programming. Vol. 248, In *Doklady akademii nauk* (5), (pp. 1049–1051). Russian Academy of Sciences.
- Kozlov, M. K., Tarasov, S. P., & Khachiyan, L. G. (1980). The polynomial solvability of convex quadratic programming. *USSR Computational Mathematics and Mathematical Physics*, 20(5), 223–228.
- Lahdelma, R., Hokkanen, J., & Salminen, P. (1998). SMAA–Stochastic Multiobjective Acceptability Analysis. *European Journal of Operational Research*, 106(1), 137–143.
- Lahdelma, R., & Salminen, P. (2001). SMAA-2: Stochastic multicriteria acceptability analysis for group decision making. *Operations Research*, 49(3), 444–454.
- Mousseau, V., Figueira, J., Dias, L., da Silva, C. G., & Climaco, J. (2003). Resolving inconsistencies among constraints on the parameters of an MCDA model. *European Journal of Operational Research*, 147(1), 72–93.
- Ojala, M., & Garriga, G. C. (2010). Permutation tests for studying classifier performance. *Journal of Machine Learning Research*, 11(6).
- Papadimitriou, C. H. (1981). On the complexity of integer programming. *Journal of the ACM*, 28(4), 765–768.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Platt, J. (1998). Sequential Minimal Optimization: A Fast Algorithm for Training Support Vector Machines. (MSR-TR-98-14), Microsoft.
- Sauré, D., & Vielma, J. P. (2019). Ellipsoidal methods for adaptive choice-based conjoint analysis. *Operations Research*, 67(2), 315–338.
- Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT Press.
- Tehrani, A. F., Stricker, M., & Hüllermeier, E. (2014). The Choquet kernel for monotone data. In *Proceedings of the Twenty-Second European Symposium on Artificial Neural Networks, ESANN 2014* (pp. 337–342).
- Van Rijsbergen, C. J. (1979). *Information retrieval* (2nd ed.). London: Butterworths.
- Wang, T., & Boutilier, C. (2003). Incremental utility elicitation with the minimax regret decision criterion. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, IJCAI 2003* (pp. 309–318).
- White, C. C., Sage, A. P., & Dozono, S. (1984). A model of multiattribute decisionmaking and trade-off weight determination under uncertainty. *IEEE Transactions on Systems, Man and Cybernetics*, (2), 223–229.
- Zaffalon, M., Corani, G., & Mauá, D. (2012). Evaluating credal classifiers by utility-discounted predictive accuracy. *International Journal of Approximate Reasoning*, 53(8), 1282–1301.
- Zhang, Z., Xu, Y., Yang, J., Li, X., & Zhang, D. (2015). A survey of sparse representation: algorithms and applications. *IEEE Access*, 3, 490–530.