

Incremental Elicitation of Rank-Dependent Aggregation Functions based on Bayesian Linear Regression

Nadjet Bourdache, Patrice Perny, Olivier Spanjaard

► **To cite this version:**

Nadjet Bourdache, Patrice Perny, Olivier Spanjaard. Incremental Elicitation of Rank-Dependent Aggregation Functions based on Bayesian Linear Regression. IJCAI-19 - Twenty-Eighth International Joint Conference on Artificial Intelligence, Aug 2019, Macao, China. pp.2023-2029. hal-02202468

HAL Id: hal-02202468

<https://hal.archives-ouvertes.fr/hal-02202468>

Submitted on 31 Jul 2019

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.

Incremental Elicitation of Rank-Dependent Aggregation Functions based on Bayesian Linear Regression

Nadjet Bourdache, Patrice Perny and Olivier Spanjaard

Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

{nadjat.bourdache, patrice.perny, olivier.spanjaard}@lip6.fr

Abstract

We introduce a new model-based incremental choice procedure for multicriteria decision support, that interleaves the analysis of the set of alternatives and the elicitation of weighting coefficients that specify the role of criteria in rank-dependent models such as ordered weighted averages (OWA) and Choquet integrals. Starting from a prior distribution on the set of weighting parameters, we propose an adaptive elicitation approach based on the minimization of the expected regret to iteratively generate preference queries. The answers of the Decision Maker are used to revise the current distribution until a solution can be recommended with sufficient confidence. We present numerical tests showing the interest of the proposed approach.

1 Introduction

Designing efficient preference elicitation methods is one of the major challenges of computer-aided decision support. Various preference models have been proposed and studied in Decision Theory for modeling individual values and generating personalized recommendations. These models use preference parameters that must be tuned to the value system of the Decision Maker (DM). In multicriteria decision support, the key parameters that must be elicited to perform the preference aggregation are weighting coefficients that specify the role of criteria in the aggregation, in particular their relative importance and sometimes also their level of interaction, see e.g., [Tehrani *et al.*, 2012]. Weighting coefficients can be elicited in a preliminary step and used for several recommendation tasks or can alternatively be elicited on the fly during the exploration of the set of alternatives. The latter approach, said to be *incremental*, has been widely investigated in AI in order to reduce the elicitation burden. Several utility elicitation methods are now available, see e.g., [Ha and Haddawy, 1997; Chajewska *et al.*, 2000; Blythe, 2002; Wang and Boutilier, 2003; Brazianas and Boutilier, 2007]. The incremental approach makes it possible to concentrate the elicitation effort on the part of the preferential information that is sufficient to determine an optimal choice. It also makes it possible to adapt the questionnaire to DM's answers, to quickly dispel the uncertainty attached to incomparable alternatives.

One key aspect of interactive preference elicitation is to limit the number of preference queries without downgrading the quality of recommendations. To this end, efficient adaptive elicitation strategies have been proposed allowing a fast reduction of the space of possible parameters until it becomes sufficiently small to be able to determine the optimal choice without ambiguity. For example, this is the case of incremental elicitation methods based on the minimization of the maximum regret proposed in [Wang and Boutilier, 2003; Boutilier *et al.*, 2006; Benabbou and Perny, 2015; Benabbou *et al.*, 2017]. Preference queries are generated to reduce the imprecision of parameters and the regrets attached to the possible decisions. The responses are translated into hard constraints limiting the space of admissible parameters and allowing sharper preferences to be derived. No opportunity is nevertheless given to the DM to contradict herself. In these approaches relying on a progressive and backtrack-free reduction of the uncertainty set, the efficiency comes at the cost of a relative lack of robustness in recommendation tasks, as illustrated by the following example.

Example 1 *We consider a set \mathcal{A} of alternatives in a bi-criteria decision problem containing 101 elements characterized by their performance vectors $\mathbf{a}^i = (i, 100 - i)$, $i = 0, \dots, 100$. We assume here that the DM's preferences can be represented by an ordered weighted average (OWA) of the form: $u_w(\mathbf{a}^i) = w \min\{a_1^i, a_2^i\} + (1 - w) \max\{a_1^i, a_2^i\}$ with $w = 0.6$. Hence, $u_{0.6}(\mathbf{a}^i) = 40 + 0.2i$, $\forall i \in \llbracket 0, 50 \rrbracket$ and $u_{0.6}(\mathbf{a}^i) = 60 - 0.2i$, $\forall i \in \llbracket 51, 100 \rrbracket$ and the $u_{0.6}$ -optimal solution is $\mathbf{a}^{50} = (50, 50)$ with utility 50. We remark here the interest of OWA aggregation that selects the better balanced vector in \mathcal{A} . Such a selection could not be performed with a linear model. Assume now that w is unknown and must be elicited. Hence the set of possible values for w is $W = [0, 1]$. Assume that the DM first declares that \mathbf{a}^k is at least as good as \mathbf{a}^i for some $i < k \leq 50$. From $u_w(\mathbf{a}^k) \geq u_w(\mathbf{a}^i)$ we obtain $w \geq 0.5$ and therefore $W = [1/2, 1]$. After this reduction, the optimal solution in \mathcal{A} is necessarily \mathbf{a}^{50} . We indeed have $u_w(\mathbf{a}^{50}) - u_w(\mathbf{a}^k) = (50 - k)(2w - 1) \geq 0$ for all $k \in \llbracket 0, 50 \rrbracket$ and $w \in [1/2, 1]$, and $u_w(\mathbf{a}^{50}) - u_w(\mathbf{a}^k) = 50 - 100w + i \geq 0$ for all $k \in \llbracket 51, 100 \rrbracket$ and $w \in [1/2, 1]$. If the DM answers that \mathbf{a}^i is strictly better than \mathbf{a}^k for some $i < k \leq 50$, a similar treatment would lead to conclude that $W = [0, 1/2]$ and that \mathbf{a}^0 and \mathbf{a}^{100} are necessarily optimal in \mathcal{A} . In both cases, an optimum is found after one preference query.*

This example shows the efficiency of incremental preference elicitation in a simple case of bi-criteria decision making using an OWA. The optimal choice can be determined while the uncertainty about criteria weights remains very large. This example also illustrates the high sensitivity of the method to possible errors made by the DM in expressing her preferences. If the DM unfortunately declares that she strictly prefers \mathbf{a}^i to \mathbf{a}^k for some $i < k \leq 50$ while its utility function is $u_{0.6}$, we may recommend \mathbf{a}^0 which is actually the worst option. This lack of robustness is due to definitive reductions of the uncertainty set W resulting from hard constraints derived from the observed preferences at every step. In case of a wrong answer, the uncertainty set will definitely exclude the correct value of the hidden parameters and reduce the relevance of recommendations. In this paper we keep unchanged the uncertainty set, and maintain a probability density function over it. The distribution is updated in a Bayesian manner.

2 Preliminaries

2.1 Rank-dependent Aggregation Functions

Let \mathcal{A} be the set of alternatives that need to be compared in order to make a decision. Any alternative $a \in \mathcal{A}$ is evaluated with respect to a set of p criteria denoted by $C = \{1, \dots, p\}$, and is characterized by a performance vector $\mathbf{a} = (a_1, \dots, a_p)$, where $a_i \in [0, 1]$ represents the utility of a with respect to criterion i . All utilities a_i are expressed on the same scale; the utility functions must be defined from the input data (criterion or attribute values), as proposed by, e.g., Grabisch and Labreuche [2010]. From now on, for simplicity, we will consider the image of \mathcal{A} in the criteria space, denoted by abuse of notation $\mathcal{A} = \{\mathbf{a}^1, \dots, \mathbf{a}^n\}$, and indifferently use the terms ‘‘alternative’’ and ‘‘performance vector’’.

A standard way to model preferences between performance vectors consists in using an aggregation function that maps vectors to scalar values. The function is called *rank-dependent* if the vector is sorted in nondecreasing order before being mapped into a scalar value. This enables to assign weighting coefficients to components of the performance vector according to their ranks. Given a vector \mathbf{a} , we denote by (\cdot) a permutation on C such that $a_{(1)} \leq \dots \leq a_{(p)}$.

We will consider two families of rank-dependent aggregation functions, namely *ordered weighted averages* (OWA) and *Choquet integrals*. These aggregation functions are known to be monotonic with respect to weak Pareto dominance (\mathbf{a} weakly dominates \mathbf{b} if $a_i \geq b_i$ for $i \in \{1, \dots, p\}$).

Ordered Weighted Averages

This family of aggregation functions has been introduced by Yager (1988). Let $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_p)$ be a weighting vector ($\sum_{i=1}^p \lambda_i = 1$ and $\lambda_i \geq 0 \forall i$). The OWA defined by $\boldsymbol{\lambda}$ reads:

$$\text{OWA}_{\boldsymbol{\lambda}}(\mathbf{a}) = \sum_{i=1}^p \lambda_i a_{(i)}. \quad (1)$$

Note that the OWA aggregator includes the min, the max and the arithmetic mean as special cases. The OWA function is concave if $\lambda_1 \geq \dots \geq \lambda_p$. The subclass of concave OWA functions is of particular interest in fair optimization because $\text{OWA}((a_1, \dots, a_i - \varepsilon, \dots, a_j + \varepsilon, \dots, a_p)) \geq \text{OWA}(\mathbf{a})$ for all i, j, ε such that $a_i - a_j \geq \varepsilon > 0$. It thus favors well-balanced performances. While instances of OWA range from the min

($\lambda_1 = 1$ and $\lambda_i = 0 \forall i > 0$) to the max ($\lambda_p = 1$ and $\lambda_i = 0 \forall i < p$) depending on the weights, instances of *concave* OWA range from the min to the arithmetic mean ($\lambda_i = 1/p \forall i$).

Example 2 Let $\mathbf{a} = (0.3, 0.5, 0.8)$ and $\mathbf{b} = (0.5, 0.6, 0.4)$. For $\boldsymbol{\lambda} = (0.5, 0.3, 0.2)$, we have $\text{OWA}_{\boldsymbol{\lambda}}(\mathbf{a}) = 0.46 < 0.47 = \text{OWA}_{\boldsymbol{\lambda}}(\mathbf{b})$, thus \mathbf{b} is preferred to \mathbf{a} .

Choquet Integrals

The Choquet integral [Schmeidler, 1986; Grabisch *et al.*, 2009] is an aggregation function that generalizes the notion of average, by defining a *capacity* on the criteria set. A capacity is a set function from 2^C to \mathbb{R} such that $v(\emptyset) = 0$, $v(C) = 1$ and $v(X) \leq v(Y)$ for all $X \subseteq Y \subseteq C$. In multicriteria analysis, $v(X)$ represents the weight attached to coalition X , for any $X \subseteq C$. The Choquet integral defined by v reads:

$$C_v(\mathbf{a}) = \sum_{i=1}^p (v(A_{(i)}) - v(A_{(i+1)})) a_{(i)} \quad (2)$$

$$= \sum_{i=1}^p [a_{(i)} - a_{(i-1)}] v(A_{(i)}) \quad (3)$$

where $A_{(i)} = \{(i), \dots, (p)\}$ (i.e., the criteria whose performance is greater or equal to $a_{(i)}$), $v(A_{(p+1)}) = 0$ and $a_{(0)} = 0$. Note that the Choquet integral includes the weighted sum as a special case: if v is additive, i.e., $v(X) = \sum_{i \in X} \lambda_i$ for some positive coefficients λ_i , then Equation 2 leads to $C_v(\mathbf{a}) = \sum_i \lambda_i a_i$. The Choquet integral also includes the OWA function as special case. Indeed, if we use a *symmetric* capacity v in Equation 2, i.e., $v(X)$ only depends on $|X|$, then $\lambda_i = v(A_{(i)}) - v(A_{(i+1)})$ only depends on i and not on the elements in $A_{(i)}$ and $A_{(i+1)}$, thus we recognize an OWA.

Example 3 Let us come back to Example 2. Consider the capacity v defined by $v(\emptyset) = 0$, $v(C) = 1$ and:

X	{1}	{2}	{3}	{1, 2}	{1, 3}	{2, 3}
$v(X)$	0.1	0.2	0.3	0.5	0.5	0.6

The evaluations of \mathbf{a} and \mathbf{b} using Equation 3 are:

$$C_v(\mathbf{a}) = 0.3 v(C) + 0.2 v(\{2, 3\}) + 0.3 v(\{3\}) = 0.51$$

$$C_v(\mathbf{b}) = 0.4 v(C) + 0.1 v(\{1, 2\}) + 0.1 v(\{2\}) = 0.47$$

Thus $C_v(\mathbf{a}) > C_v(\mathbf{b})$ and \mathbf{a} is preferred to \mathbf{b} .

We now recall a useful representation of capacities in terms of *Möbius inverse*. The Möbius inverse of a capacity v is a mapping $m_v : 2^C \rightarrow \mathbb{R}$ defined by $m_v(X) = \sum_{Y \subseteq X} (-1)^{|X \setminus Y|} v(Y) \forall X \subseteq C$. The quantity $m_v(X)$ is called the *Möbius mass* of X . The capacity can be derived from its Möbius inverse as follows: $\forall X \subseteq C$, $v(X) = \sum_{Y \subseteq X} m_v(Y)$. The Möbius inverse yields an alternative formulation of the Choquet integral:

$$C_v(\mathbf{a}) = \sum_{X \subseteq C} m_v(X) \min_{i \in X} a_i$$

A capacity v is said to be *k-additive* if $m_v(X) = 0$ for all $X \subseteq C$ such that $|X| > k$ and there exists at least one subset X such that $|X| = k$ and $m_v(X) > 0$. For instance, the capacity v considered in Example 3 is at most 2-additive because: $m(\{1, 2, 3\}) = v(\{1, 2, 3\}) - \sum_{i \neq j} v(\{i, j\}) + \sum_i v(\{i\}) = 0$. The main interest of Choquet integrals for multicriteria decision making stems from the fact that they can represent interaction between criteria [Grabisch and Labreuche, 2010], as can be witnessed by the Möbius masses. For instance, by definition, $m_v(\{i, j\}) = v(\{i, j\}) - v(\{i\}) - v(\{j\})$. There is a positive interaction between i and j if $m_v(\{i, j\}) > 0$,

and negative if $m_v(\{i, j\}) < 0$. The k -additive capacities for small values of k greater than 1 are very useful because in practical situations, they offer a sufficient expressivity to model positive or negative synergies between criteria with a reduced number of parameters. For example, when $k = 2$ the capacity is completely characterized by $(p^2 + p)/2$ coefficients (one Möbius mass for every singleton and every pair).

The aggregation functions introduced above are intended to be used to describe the preferences of the DM, thus their parameters (i.e., the weighting vector or the capacity) must be inferred from preference statements provided by the DM.

2.2 Related Works

Regarding the Bayesian incremental elicitation of an aggregation function from pairwise preference queries, the previous works in the literature can be divided in two categories: non-parametric approaches and parametric approaches.

Following a *nonparametric approach*, Chu and Ghahramani [2005] proposed to use Gaussian Processes (GP) for preference learning. To approximate *any* possible aggregation function, the idea is to assume that each overall utility value follows a normal distribution. The joint distribution defines a GP. This distribution is characterized by a mean function and a covariance matrix, and can be seen as an infinite-dimensional extension of a multivariate Gaussian distribution (the domain of the aggregation function can indeed be infinite). The covariance matrix is used to ensure that alternatives with similar performance vectors yield similar overall utility values. Building on this work, Brochu *et al.* [2007] proposed an active learning procedure that does not require to accurately model the entire aggregation function to determine an optimal alternative. Their approach is based on pairwise comparison queries. Zintgraf *et al.* [2018] recently extended it to other query types, like asking the user to rank or cluster alternatives. In this paper, we investigate a parametric approach where one imposes a rank-dependent aggregation function to ensure appealing normative properties.

A *parametric approach* consists indeed in learning the parameters of a given aggregation function. This kind of approaches was well introduced and illustrated in [Chajewska *et al.*, 2000] for the adaptive elicitation of utility values over a finite set of outcomes in decision making under uncertainty. While their approach used standard gamble queries (pairwise preference queries between probability distributions over outcomes), Guo and Sanner [2010] proposed a variant that only requires the users to state their preference under certainty, so as to reduce the cognitive load. Recently, for the adaptive elicitation of the parameters of a multiattribute utility function, Sauré and Vielma [2018] proposed an approach that updates an ellipsoidal credibility region computed from a multivariate normal distribution over the space of parameters, and showed how to use mixed integer programming to determine queries that are likely to reduce the volume of the credibility region. Nonetheless, none of these works is concerned with learning the parameters of a *rank-dependent* aggregation function, which is the aim of the present paper.

Most of the works aiming at determining the parameters of rank-dependent aggregation functions consider a static preference database as input, and focus on the determination of

parameter values that best fit the available preferences. For learning the parameters of a Choquet integral, Marichal and Roubens [2000] presented a linear programming formulation, Tehrani *et al.* [2012] proposed a generalization of logistic regression, and Rowley *et al.* [2015] used a principal component analysis. Departing from these works, incremental elicitation procedures have recently been proposed [Benabbou *et al.*, 2017; Bourdache and Perny, 2019], but they assume no response-error. Yet, as illustrated in Example 1, the recommendation made can be severely impacted by errors in responses. To overcome this drawback, we introduce in the next section some useful linear and compact reformulations of OWA and Choquet aggregators in new multidimensional spaces, that makes it possible to implement Bayesian linear regression techniques with rank-dependent models.

3 Basis Functions for OWA and Choquet

A way to extend Bayesian linear regression to a nonlinear function f is to reformulate f as a linear combination of (non-linear) *basis* functions g_1, \dots, g_q [Bishop, 2006]. Formally:

$$f(\mathbf{a}) = \sum_{i=1}^q w_i g_i(\mathbf{a}) = \mathbf{w}^T \mathbf{g}(\mathbf{a})$$

where $\mathbf{w} = (w_1, \dots, w_q)^T$ is a weighting vector and $\mathbf{g}(\mathbf{a}) = (g_1(\mathbf{a}), \dots, g_q(\mathbf{a}))^T$. Thus, f is a linear function in the space of features g_1, \dots, g_q . We now introduce suitable basis functions for the rank-dependent aggregators studied here.

Ordered Weighted Averages

Obviously, any OWA function is a linear combination of functions $g_i(\mathbf{a}) = a_{(i)}$ for $i = 1, \dots, p$ (see Equation 1). Given a weighting vector λ for OWA, the corresponding parameters for the basis functions are $w_i = \lambda_i$ for $i = 1, \dots, p$. For the OWA function to be nondecreasing in every component, we impose $w_i \geq 0$ for $i = 1, \dots, p$.

Concave Ordered Weighted Averages

In the context of fair optimization, the constraint $w_1 \geq \dots \geq w_p$ is required to favor solutions having balanced profile [Weymark, 1981], which makes the linearization more tricky. In this case we get a concave OWA function. To overcome the problem, we shall use cumulative functions defined by $g_i(\mathbf{a}) = \sum_{k=1}^i a_{(k)}$. It can easily be checked that:

$$\text{OWA}_\lambda(\mathbf{a}) = \sum_{i=1}^p w_i \left(\sum_{k=1}^i a_{(k)} \right)$$

where $w_i = \lambda_i - \lambda_{i+1}$ for $i < p$ and $w_p = \lambda_p$. Conversely, given parameters w_k , the corresponding weights for OWA are $\lambda_i = \sum_{k=i}^p w_k$. Note that to enforce constraint $\lambda_1 \geq \dots \geq \lambda_p \geq 0$, we only have to impose $w_i \geq 0$ for $i = 1, \dots, p$.

Example 4 For $\lambda = (0.5, 0.3, 0.2)$ (as in Example 2), the expression $\text{OWA}_\lambda(\mathbf{a})$ reformulates as $0.2g_1(\mathbf{a}) + 0.1g_2(\mathbf{a}) + 0.2g_3(\mathbf{a})$. If $\mathbf{a} = (0.3, 0.5, 0.8)$, we have $\mathbf{g}(\mathbf{a}) = (0.3, 0.8, 1.6)$ and $\text{OWA}_\lambda(\mathbf{a}) = 0.2 \times 0.3 + 0.1 \times 0.8 + 0.2 \times 1.6 = 0.46$.

2-additive Choquet Integrals

To express a Choquet integral as a linear combination of basis functions, we consider a decomposition of v under the form $v = \sum_{i=1}^q w_i v_i$, where v_i 's are basis capacity functions. Indeed, if $v = \sum_{i=1}^q w_i v_i$, then $C_v(\mathbf{a}) = \sum_{i=1}^q w_i C_{v_i}(\mathbf{a})$ because C_v is linear in v . For the class of at most 2-additive capacities,

a decomposition of v with $q = p^2$ basis functions is known. This class can be characterized by a polytope, the extreme points of which are defined by two families of 0-1 capacities:

- the *unanimity games* for singletons and pairs defined by

$$v_i(X) = \begin{cases} 1 & \text{if } Y_i \subseteq X \\ 0 & \text{otherwise} \end{cases} \quad \text{for } i \in \llbracket 1, \frac{p(p+1)}{2} \rrbracket$$

where $Y_i \subseteq C$ is any nonempty subset of size at most 2;

- the *conjugates* of unanimity games for pairs defined by

$$v_i(X) = \begin{cases} 1 & \text{if } Y_i \cap X \neq \emptyset \\ 0 & \text{otherwise} \end{cases} \quad \text{for } i \in \llbracket \frac{p(p+1)}{2} + 1, p^2 \rrbracket$$

where $Y_i \subseteq C$ is any subset of size 2.

This characterization of extreme points directly follows from [Grabisch, 2016, Theorem 2.65, p.89]. Hence, for any at most 2-additive capacity, there exists p^2 coefficients $w_i \geq 0$ adding up to 1 such that $C_v(\mathbf{a}) = \sum_i w_i C_{v_i}(\mathbf{a})$. Note that a similar result, formulated in terms of Möbius inverse, has also been used in [Hüllermeier and Tehrani, 2013].

Example 5 Any capacity v on three criteria can be written as a convex combination of v_i 's in the table below (the unanimity games are v_1, \dots, v_6 , the conjugates are v_7, v_8, v_9).

X	{1}	{2}	{3}	{1,2}	{1,3}	{2,3}
$v_1(X)$	1	0	0	1	1	0
$v_2(X)$	0	1	0	1	0	1
$v_3(X)$	0	0	1	0	1	1
$v_4(X)$	0	0	0	1	0	0
$v_5(X)$	0	0	0	0	1	0
$v_6(X)$	0	0	0	0	0	1
$v_7(X)$	1	1	0	1	1	1
$v_8(X)$	1	0	1	1	1	1
$v_9(X)$	0	1	1	1	1	1

For $\mathbf{a} = (0.3, 0.5, 0.8)$, we have $C_{v_1}(\mathbf{a}) = 0.3$, $C_{v_2}(\mathbf{a}) = 0.5$, $C_{v_3}(\mathbf{a}) = 0.8$, $C_{v_4}(\mathbf{a}) = 0.3$, $C_{v_5}(\mathbf{a}) = 0.3$, $C_{v_6}(\mathbf{a}) = 0.5$, $C_{v_7}(\mathbf{a}) = 0.5$, $C_{v_8}(\mathbf{a}) = 0.8$, $C_{v_9}(\mathbf{a}) = 0.8$. Besides, $v = 0.1v_2 + 0.3v_3 + 0.3v_4 + 0.1v_5 + 0.1v_6 + 0.1v_7$, thus $C_v(\mathbf{a}) = 0.51$.

4 Incremental Elicitation Process

We now introduce an incremental Bayesian elicitation procedure taking advantage of the above linear reformulations. We consider a non-linear aggregation function admitting a linear representation $f_{\mathbf{w}}(\mathbf{a}) = \mathbf{w}^T \mathbf{g}(\mathbf{a})$, where values $g_i(\mathbf{a})$ are computed as shown in Section 3. Incremental decision making consists in implementing an adaptive elicitation procedure based on $f_{\mathbf{w}}$ to iteratively collect preference statements so as to reduce the uncertainty about \mathbf{w} until an alternative can be recommended with sufficient confidence. We propose a Bayesian approach where a prior density function over the parameter space is iteratively updated with new preference statements. The incremental decision process is based on the progressive minimization of expected regrets attached to possible decisions. The expected regrets are defined as follows.

Definition 1 Given a density function p on the possible values of \mathbf{w} , the pairwise expected regret of an alternative \mathbf{a} with respect to an alternative \mathbf{b} is defined as follows:

$$\text{PER}(\mathbf{a}, \mathbf{b}, p) = \int \max\{0, f_{\mathbf{w}}(\mathbf{b}) - f_{\mathbf{w}}(\mathbf{a})\} p(\mathbf{w}) d\mathbf{w}$$

In other words, the pairwise expected regret of \mathbf{a} with respect to \mathbf{b} represents the expected utility loss when recommending \mathbf{a} instead of \mathbf{b} .

Definition 2 Given a set \mathcal{A} of alternatives and a density function p on the possible values of \mathbf{w} , the max expected regret of $\mathbf{a} \in \mathcal{A}$ and the minimax expected regret are defined by:

$$\text{MER}(\mathbf{a}, \mathcal{A}, p) = \max_{\mathbf{b} \in \mathcal{A}} \text{PER}(\mathbf{a}, \mathbf{b}, p)$$

$$\text{MMER}(\mathcal{A}, p) = \min_{\mathbf{a} \in \mathcal{A}} \text{MER}(\mathbf{a}, \mathcal{A}, p)$$

In other words, the max expected regret¹ of \mathbf{a} is the maximum expected utility loss incurred in selecting \mathbf{a} in \mathcal{A} .

Incremental Elicitation

We denote by $\mathbf{a}^{(k)}$ an alternative achieving the $\text{MMER}(\mathcal{A}, p)$ value at iteration k . This alternative is worth recommending if no further preferential information can be obtained because it minimizes the expected opportunity loss. Otherwise, we ask the DM to compare $\mathbf{a}^{(k)}$ and any adversarial choice $\mathbf{b}^{(k)}$ in $\arg \max_{\mathbf{b} \in \mathcal{A}} \text{PER}(\mathbf{a}^{(k)}, \mathbf{b}, p)$, as suggested by Boutilier *et al.* (2006) in the context of max regret minimization (*current solution strategy*). The DM's answer allows us to update the current distribution p for the next iteration. This process can be iterated until the MMER value drops below a given acceptance threshold δ . The resulting incremental decision procedure is more formally described in Algorithm 1.

Algorithm 1 Incremental decision making

Input: \mathcal{A} : alternatives; δ : threshold; $p(\mathbf{w})$: density function.

Output: \mathbf{a}^* : best recommendation in \mathcal{A} .

```

1  $i \leftarrow 1$ ;
2 repeat
3    $\mathbf{a}^{(i)} \leftarrow \arg \min_{\mathbf{a} \in \mathcal{A}} \text{MER}(\mathbf{a}, \mathcal{A}, p)$ 
4    $\mathbf{b}^{(i)} \leftarrow \arg \max_{\mathbf{b} \in \mathcal{A}} \text{PER}(\mathbf{a}^{(i)}, \mathbf{b}, p)$ 
5   Ask the DM if  $\mathbf{a}^{(i)}$  is preferred to  $\mathbf{b}^{(i)}$ 
6    $y^{(i)} \leftarrow 1$  if the answer is yes and 0 otherwise
7    $p(\mathbf{w}) \leftarrow p(\mathbf{w}|y^{(i)}) \quad \forall \mathbf{w}$  (see Algorithm 2)
8    $i \leftarrow i + 1$ 
9 until  $\text{MMER}(\mathcal{A}, p) \leq \delta$ 
10 return  $\mathbf{a}^*$  selected in  $\arg \min_{\mathbf{a} \in \mathcal{A}} \text{MER}(\mathbf{a}, \mathcal{A}, p)$ 

```

Bayesian Inference

In Line 7, the density function $p(\mathbf{w})$ is updated using the latent variable approach introduced by Albert and Chib (1993) for binary *probit* regression. Let $y^{(i)}$ represent the answer to a comparison query i between $\mathbf{a}^{(i)}$ and $\mathbf{b}^{(i)}$, with $y^{(i)} = 1$ if $\mathbf{a}^{(i)}$ is preferred, $y^{(i)} = 0$ otherwise. In our setting, the probit model is used to predict the value of $y^{(i)}$ based on the vector $\mathbf{d}^{(i)} = \mathbf{g}(\mathbf{a}^{(i)}) - \mathbf{g}(\mathbf{b}^{(i)})$ of explanatory variables. The answer $y^{(i)}$ is assumed to depend on a latent variable $z^{(i)}$ representing preference intensity and defined by:

$$z^{(i)} = \mathbf{w}^T \mathbf{d}^{(i)} + \varepsilon^{(i)} \quad (4)$$

where $\varepsilon^{(i)} \sim \mathcal{N}(0, 1)$ is a Gaussian noise accounting for possible errors in the answers and for descriptive limitations of the preference model. If $z^{(i)} \geq 0$, then $y^{(i)} = 1$, else $y^{(i)} = 0$.

Following the sequential Bayesian approach of Algorithm 1, a prior density function $p(\mathbf{w})$ is defined on the parameter space. In the probit model, $p(\mathbf{w})$ is multivariate Gaussian. After each query, the current density function $p(\mathbf{w})$

¹Note that there exists similar definitions of expected regrets in the literature (see, e.g., Guo and Sanner [2010]).

is revised according to answer $y^{(i)}$. The main interest of resorting to the probit model is to obtain a computationally efficient procedure to estimate the posterior $p(\mathbf{w}|y^{(i)})$. The revised density function $p(\mathbf{w}|y^{(i)})$ can indeed be formulated as a marginal density of $p(\mathbf{w}, z|y^{(i)})$:

$$p(\mathbf{w}|y^{(i)}) = \int p(\mathbf{w}, z|y^{(i)}) dz = \int p(\mathbf{w}|z)p(z|y^{(i)}) dz \quad (5)$$

However, $p(z|y^{(i)})$ depends, in turn, on $p(\mathbf{w}|y^{(i)})$ because:

$$p(z|y^{(i)}) = \int p(\mathbf{w}|y^{(i)})p(z|\mathbf{w}, y^{(i)}) d\mathbf{w} \quad (6)$$

We use an iterative method to approximate the density function $p(\mathbf{w}|y^{(i)})$ that solves the fixed point equation obtained by replacing $p(z|y^{(i)})$ by (6) in Equation (5), as suggested by Tanner and Wong (1987). They proved that it should converge to the true posterior $p(\mathbf{w}|y^{(i)})$. The method is detailed in Algorithm 2, where the upper bound K on the number of iterations depends on the compromise sought between speed and accuracy. At each iteration k , an approximation $p_k(\mathbf{w})$ of $p(\mathbf{w}|y^{(i)})$ is refined. The for loop in Line 3 aims at sampling m values z_1, \dots, z_m from $p(z|y^{(i)})$. The density function $p(z|\mathbf{w}, y^{(i)})$ used in line 5 is truncated Gaussian:

$$p(z|\mathbf{w}, y^{(i)}) \propto \begin{cases} \mathcal{N}(\omega^T \mathbf{d}^{(i)}, I) \mathbb{1}_{z \geq 0} & \text{if } y^{(i)} = 1 \\ \mathcal{N}(\omega^T \mathbf{d}^{(i)}, I) \mathbb{1}_{z < 0} & \text{otherwise} \end{cases}$$

From sampled values z_j , a tighter approximation $p_{k+1}(\mathbf{w})$ is computed in Line 6. Note that $\frac{1}{m} \sum_{j=1}^m p_k(\mathbf{w}|z_j)$ is multivariate Gaussian because each $p_k(\mathbf{w}|z_j)$ is multivariate Gaussian. It is indeed well-known (see e.g. Bishop, 2006) that if $p_k(\mathbf{w}) = \mathcal{N}(\boldsymbol{\mu}_k, \mathbf{S}_k)$, then $p_k(\mathbf{w}|z_j) = \mathcal{N}(\boldsymbol{\mu}^j, \mathbf{S})$ where:

$$\mathbf{S}^{-1} = \mathbf{d}^{(i)T} \mathbf{d}^{(i)} + \mathbf{S}_k^{-1} \text{ and } \boldsymbol{\mu}^j = \mathbf{S}(\mathbf{d}^{(i)T} z_j + \mathbf{S}_k^{-1} \boldsymbol{\mu}_k)$$

Thus $\frac{1}{m} \sum_j p_k(\mathbf{w}|z_j) = \frac{1}{m} \sum_j \mathcal{N}(\boldsymbol{\mu}^j, \mathbf{S}) = \mathcal{N}(\frac{1}{m} \sum_j \boldsymbol{\mu}^j, \mathbf{S})$.

Coming back to the interpretation of \mathbf{w} as weights in a linear reformulation $f_{\mathbf{w}}$, the conditions $\sum_i \lambda_i = 1$ (for OWA) and $v(C) = 1$ (for Choquet) turn into normalisation constraints $\sum_i i \cdot w_i = 1$ (for concave OWA) and $\sum_i w_i = 1$ (for general OWA and Choquet). As functions $f_{\mathbf{w}}$ are linear in \mathbf{w} (thus the induced preferences are invariant by rescaling of \mathbf{w}), the normalisation constraints can be dropped, which considerably simplifies the sampling. In the calculation of regrets in Algorithm 1, we use nonetheless normalised weights to compare regrets from one iteration to another. Besides, we use a

Algorithm 2 Approximating density function $p(\mathbf{w}|y^{(i)})$

Input: $p(\mathbf{w})$: density function; m : sample size; K : #iterations.
Output: $p(\mathbf{w}|y^{(i)})$: posterior density function.
1 $k \leftarrow 0$; $p_0(\mathbf{w}) \leftarrow p(\mathbf{w}) \quad \forall \mathbf{w}$
2 **repeat** /* $p_k(\mathbf{w})$ approximates $p(\mathbf{w}|y^{(i)})$ */
3 **for** $j=1$ to m **do**
4 draw ω from $p_k(\mathbf{w})$
5 draw z_j from $p(z|\omega, y^{(i)})$
6 $p_{k+1}(\mathbf{w}) \leftarrow \frac{1}{m} \sum_{j=1}^m p_k(\mathbf{w}|z_j) \quad \forall \mathbf{w}$
7 $k \leftarrow k + 1$
8 **until** p_k stabilizes **or** $k=K$
9 **return** p_k

truncated normal density function $p(\mathbf{w})$ to ensure $w_i \geq 0 \quad \forall i$.

5 Numerical Tests

Before coming to the results of numerical tests² we conducted on randomly generated instances to evaluate the behavior of Algorithm 1, we illustrate this behavior on Example 1.

Example 1 (continued). *Assume that Algorithm 1 is launched with an acceptance threshold $\delta = 0.02$. We first obtain $\text{MMER} = 0.05 > \delta$. The DM is then asked to compare \mathbf{a}^{24} and \mathbf{a}^{50} (this is not necessarily \mathbf{a}^{50} , it may be a close neighbor depending on the sample drawn for w). Assume that the DM makes a mistake and declares that she prefers \mathbf{a}^{24} to \mathbf{a}^{50} . Using this information the probability density function over weights is revised and MMER is recomputed. Now, $\text{MMER} = 0.04 > \delta$ so a second query is generated. The DM compares \mathbf{a}^{11} and \mathbf{a}^{50} and correctly declares that she prefers \mathbf{a}^{50} . We obtain $\text{MMER} = 0.05 > \delta$, and afterwards four additional queries are needed to conclude. The DM is asked to compare \mathbf{a}^{28} , \mathbf{a}^{34} , \mathbf{a}^{43} and \mathbf{a}^{46} to \mathbf{a}^0 respectively and the DM never chooses \mathbf{a}^0 . After the six answers, we obtain $\text{MMER} = 0.01 < \delta$ and the algorithm stops recommending $\arg \min_{\mathbf{a} \in \mathcal{A}} \text{MER}(\mathbf{a}, \mathcal{A}, p) = \mathbf{a}^{47}$, which is very close to the actual optimal solution \mathbf{a}^{50} .*

We thus observe that \mathbf{a}^{47} is now recommended while \mathbf{a}^0 or \mathbf{a}^{100} was recommended using a deterministic approach with the same wrong initial answer. For the hidden utility $u_{0.6}$ of the DM, we have $u_{0.6}(\mathbf{a}^{50}) = 40 + 0.2 \times 50 = 50$ and $u_{0.6}(\mathbf{a}^{47}) = 40 + 0.2 \times 47 = 49.4$. The actual regret in recommending \mathbf{a}^{47} instead of \mathbf{a}^{50} is therefore $50 - 49.4 = 0.6$, which represents 1.2% of the utility value of an optimal alternative.

Random Generation of Instances

We generate instances with 5 criteria and 100 Pareto optimal alternatives. Every alternative \mathbf{a} in each generated set \mathcal{A} is drawn as follows: a first vector \mathbf{v} of size $p-1$ is uniformly drawn in $[0, 1]^{p-1}$, then \mathbf{a} is obtained by setting $a_i = v_{(i)} - v_{(i-1)}$ for $i=1, \dots, p$, where $v_{(0)} = 0$ and $v_{(p)} = 1$. To avoid that all generated alternatives share the same hyperplane $\sum_i a_i = 1$ in the utility space, we apply the square root function on all components a_i of each performance vector $\mathbf{a} \in \mathcal{A}$. Finally, for each instance, a hidden weight vector \mathbf{w} is uniformly drawn over the simplex. For the elicitation of OWA parameters, the prior is set to $\mathcal{N}((10, \dots, 10), 100\mathbf{I})$, where \mathbf{I} is the identity matrix. For Choquet parameters, the prior is set to $\mathcal{N}(\boldsymbol{\mu}, 100\mathbf{I})$, where $\mu_i = 10$ if $|Y_i| = 1$ and $\mu_i = 0$ otherwise (we recall that $C_v(\mathbf{a}) = \sum_i w_i C_{v_i}(\mathbf{a})$, where each v_i is attached to a subset Y_i of size at most 2). In this way, the prior density function is centered around a parameter vector for which there is no interaction between criteria.

Simulation of the Interactions with the DM

To simulate the DM's answers, we use the model given in Equation 4 with $\varepsilon^{(i)} \sim \mathcal{N}(0, \sigma^2)$ for $\sigma \in \{0, 0.1, 0.2\}$, in order to test the tolerance of the approach to wrong answers (none

²Implementation in Python using the `tmvtnorm` R's library to draw vectors according to multivariate truncated normal densities, running on an Intel(R) Core(TM) i7-4790 CPU with 15GB of RAM.

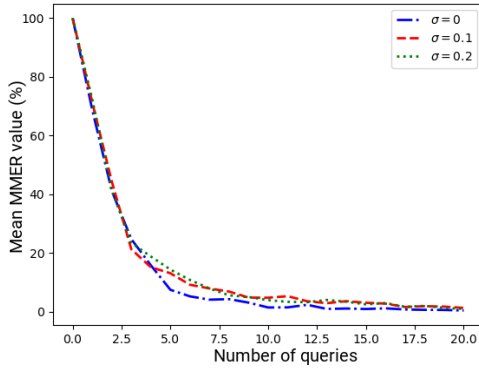


Figure 1: Mean MMER value (OWA)

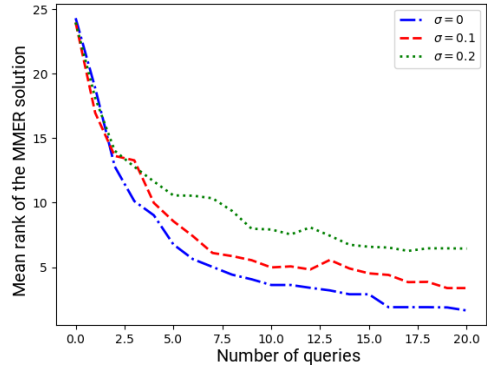


Figure 2: Mean rank of the recommendation (OWA)

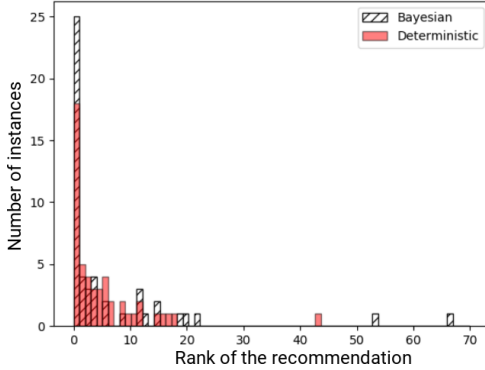


Figure 3: Rank of the recommendation (OWA)

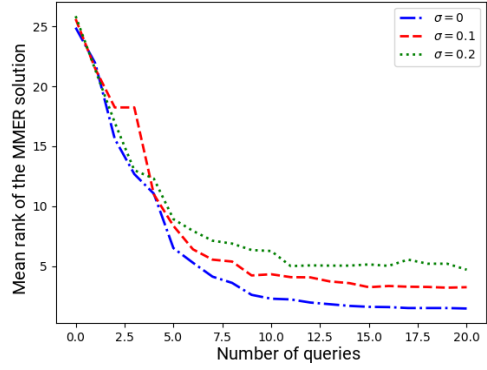


Figure 4: Mean rank of the recommendation (Choquet)

for $\sigma = 0$). Setting $\sigma = 0.1$ (resp. 0.2) led to 19% (resp. 26%) of wrong answers on the considered runs of the algorithm for OWA, and 10% (resp. 17%) for 2-additive Choquet.

Analysis of the Results

We evaluated the efficiency of Algorithm 1 with respect to the value of σ (results averaged over 50 instances).

We first observed the evolution of the MMER value in the course of Algorithm 1, in percentage of the initial MMER value (the curves for OWA and 2-additive Choquet are very similar, thus only the former are provided), as well as the evolution of the actual rank in \mathcal{A} of the current MMER alternative according to the hidden DM preference model. The obtained curves for OWA and at most 2-additive Choquet integrals are represented in Figure 1, 2 and 4. We observe that the mean MMER value drops below 3% of the initial MMER value in a dozen queries. The value of σ mainly impacts the quality of the recommendation, measured by the rank of the recommended alternative in the hidden ranking of the DM, that deteriorates when the value of σ increases. Nevertheless, the average rank of the recommended alternative is never beyond 5 (over 100 alternatives), except for OWA with $\sigma = 0.2$ where it is around 6.5. Regarding the computation times between two queries, it takes about 5 seconds for OWA, and about 35 seconds for 2-additive Choquet. This difference can be explained by the respective sizes of parameter spaces.

We next compared the performance of Algorithm 1 to a deterministic approach based on the minimization of the max regret, that reduces the parameter space after each query instead

of updating a probability density function (see e.g., [Bourdache and Perny, 2019]). For both methods, we considered two performance criteria: the number of queries for making a recommendation, and the quality of the recommendation itself, measured in terms of actual rank of the recommended solution at termination. Regarding the number of queries, as can be seen in Figure 2 and 4, the quality of the recommendation does not significantly improve after a dozen query, which is similar to the number of queries in the deterministic approach. Regarding the quality of the recommendation, the histogram in Figure 3 shows the distribution of the ranks of the recommendations over the 50 randomly generated instances, for OWA with $\sigma = 0.2$. The Bayesian approach recommends the best alternative in 50% of cases while it happens in 36% of cases for the deterministic approach.

6 Conclusion

Our work differentiates from previous works on incremental decision making with OWA and Choquet integrals in proposing a probabilistic approach that is tolerant to response-errors of the DM. It also differentiates from previous work on active preference learning with Bayesian methods in implementing an approach that applies to well-known rank-dependent decision criteria. Our contribution on the Choquet integral concerns the case of 2-additive capacities, that are known to provide sufficient flexibility to model very diverse preference systems including positive and negative interactions between criteria. However, it would be interesting to extend our approach to elicit larger classes of capacities to make even bet-

ter use of the descriptive possibilities offered by the Choquet integral. This is a challenging problem because the polytope of at most k -additive capacities for $k > 2$ admits many other extreme points than 0-1 capacities [Miranda *et al.*, 2006].

References

- [Albert and Chib, 1993] James H Albert and Siddhartha Chib. Bayesian analysis of binary and polychotomous response data. *J. Am. Stat. Assoc.*, 88(422):669–679, 1993.
- [Benabbou and Perny, 2015] Nawal Benabbou and Patrice Perny. Incremental weight elicitation for multiobjective state space search. In *AAAI-15*, pages 1093–1099, 2015.
- [Benabbou *et al.*, 2017] Nawal Benabbou, Patrice Perny, and Paolo Viappiani. Incremental elicitation of Choquet capacities for multicriteria choice, ranking and sorting problems. *Artificial Intelligence*, 246:152–180, 2017.
- [Bishop, 2006] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [Blythe, 2002] Jim Blythe. Visual exploration and incremental utility elicitation. In *AAAI-02*, pages 526–532, 2002.
- [Bourdache and Perny, 2019] Nadjet Bourdache and Patrice Perny. Active preference elicitation based on generalized gini functions: Application to the multiagent knapsack problem. In *AAAI 2019*, 2019.
- [Boutilier *et al.*, 2006] Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 170(8-9):686–713, 2006.
- [Braziunas and Boutilier, 2007] Darius Braziunas and Craig Boutilier. Minimax regret based elicitation of generalized additive utilities. In *Proc. of UAI-07*, pages 25–32, 2007.
- [Brochu *et al.*, 2007] Eric Brochu, Nando D Freitas, and Abhijeet Ghosh. Active preference learning with discrete choice data. In *NIPS 2007*, pages 409–416, 2007.
- [Chajewska *et al.*, 2000] Urszula Chajewska, Daphne Koller, and Ronald Parr. Making rational decisions using adaptive utility elicitation. In *Proceedings of AAAI-00*, pages 363–369, 2000.
- [Chu and Ghahramani, 2005] Wei Chu and Zoubin Ghahramani. Preference learning with gaussian processes. In *ICML-05*, pages 137–144, 2005.
- [Grabisch and Labreuche, 2010] Michel Grabisch and Christophe Labreuche. A decade of application of the Choquet and Sugeno integrals in multi-criteria decision aid. *Annals of OR*, 175(1):247–286, 2010.
- [Grabisch *et al.*, 2009] Michel Grabisch, Jean-Luc Marichal, Radko Mesiar, and Endre Pap. *Aggregation Functions*. Camb. U. Press, 2009.
- [Grabisch, 2016] Michel Grabisch. *Set functions, games and capacities in decision making*, volume 46. Springer, 2016.
- [Guo and Sanner, 2010] Shengbo Guo and Scott Sanner. Multiattribute bayesian preference elicitation with pairwise comparison queries. In *NIPS*, pages 396–403, 2010.
- [Ha and Haddawy, 1997] Vu Ha and Peter Haddawy. Problem-focused incremental elicitation of multi-attribute utility models. In *Proc. of UAI-97*, pages 215–222, 1997.
- [Hüllermeier and Tehrani, 2013] Eyke Hüllermeier and Ali Fallah Tehrani. Efficient learning of classifiers based on the 2-additive Choquet integral. In *Comp. Intelligence in Intelligent Data Analysis*, pages 17–29. Springer, 2013.
- [Marichal and Roubens, 2000] Jean-Luc Marichal and Marc Roubens. Determination of weights of interacting criteria from a reference set. *European Journal of Operational Research*, 124(3):641–650, 2000.
- [Miranda *et al.*, 2006] Pedro Miranda, Elías F. Combarro, and Pedro Gil. Extreme points of some families of non-additive measures. *European Journal of Operational Research*, 174(3):1865–1884, 2006.
- [Rowley *et al.*, 2015] Hazel V Rowley, Arne Geschke, and Manfred Lenzen. A practical approach for estimating weights of interacting criteria from profile sets. *Fuzzy Sets and Systems*, 272:70–88, 2015.
- [Sauré and Vielma, 2018] Denis Sauré and Juan Pablo Vielma. Ellipsoidal methods for adaptive choice-based conjoint analysis. *Operations Research*, 2018.
- [Schmeidler, 1986] David Schmeidler. Integral representation without additivity. *Proceedings of the American mathematical society*, 97(2):255–261, 1986.
- [Tanner and Wong, 1987] Martin A Tanner and Wing Hung Wong. The calculation of posterior distributions by data augmentation. *J. Am. Stat. Assoc.*, 82(398):528–540, 1987.
- [Tehrani *et al.*, 2012] Ali Fallah Tehrani, Weiwei Cheng, Krzysztof Dembczynski, and Eyke Hüllermeier. Learning monotone nonlinear models using the Choquet integral. *Machine Learning*, 89(1-2):183–211, 2012.
- [Wang and Boutilier, 2003] Tianhan Wang and Craig Boutilier. Incremental utility elicitation with the minimax regret decision criterion. In *IJCAI-03*, pages 309–318, 2003.
- [Weymark, 1981] John A Weymark. Generalized gini inequality indices. *Mathematical Social Sciences*, 1(4):409–430, 1981.
- [Yager, 1988] Ronald R Yager. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Transactions on Systems, Man, and Cybernetics*, 18(1):183–190, 1988.
- [Zintgraf *et al.*, 2018] Luisa M Zintgraf, Diederik M Roijers, Sjoerd Linders, Catholijn M Jonker, and Ann Nowé. Ordered preference elicitation strategies for supporting multi-objective decision making. In *Proceedings of AAMAS-18*, pages 1477–1485, 2018.