Adaptive Elicitation of Rank-Dependent Aggregation Models based on Bayesian Linear Regression

Nadjet Bourdache[1] and **Patrice Perny**[1] and **Olivier Spanjaard**[1]

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 to capture the value system of the Decision Maker (DM) is one of the major challenges of computer-aided decision support. In the field of 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., [10]. 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 and OR in order to reduce the elicitation burden, see e.g., [13, 9, 8, 4, 12, 2].

One key aspect of interactive preference elicitation is to limit the number of preference queries without downgrading the quality of recommendations. For example, this is the goal of incremental elicitation methods based on the minimization of the maximum regret proposed in [12, 3] where the answers to queries are translated into hard constraints limiting the space of admissible parameters and allowing to refine current preferences. No opportunity is nevertheless given to the DM to contradict itself, thus the efficiency comes at the cost of a relative lack of robustness in recommendation tasks.

A way to be less sensitive to errors is to adopt a probabilistic approach. Starting from a prior distribution over a class of admissible preference models, the observation of new preference statements can be used to incrementally revise the current model and progressively reduce the risk attached to recommendations. This revision can be embodied into an adaptive elicitation process where preference queries are selected according to their expected value of information. This approach is well introduced and illustrated in [4] for the elicitation of utilities. Multiple variants relying on the Bayesian approach have been proposed, see e.g., [5, 7, 11]. Following this line, our aim in this work is to propose a Bayesian approach for the incremental elicitation of weighting parameters in *rank-dependent* (and therefore non-linear) multicriteria decision models such as ordered weighted averages [14] and Choquet integrals [6], where preference queries are selected by *expected regret minimization*. The goal is not to assess the weighting parameters of these models precisely but only to reduce the uncertainty sufficiently to be able to make a choice within a given set of alternatives, with the desired level of confidence.

2 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 $C = \{1, \ldots, p\}$, and is characterized by a performance vector $\mathbf{a} = (a_1, \ldots, a_p)$, where $a_i \in [0, 1]$ represents the utility of a with respect to criterion i. From now on, for simplicity, we will consider the image of $\mathcal{A}$ in the criteria space, denoted by abuse of notation $\mathcal{A} = \{\mathbf{a^1}, \ldots, \mathbf{a^n}\}$, and indifferently use the terms "alternative" and "performance vector". Given a performance vector $\mathbf{a}$, we denote by $(\cdot)$ a permutation on C such that $a_{(1)} \leq \ldots \leq a_{(p)}$. Thus, $a_{(i)}$ represents the smallest i^{th} component of $\mathbf{a}$. We will consider two families of rank-dependent aggregation functions:

– The *ordered weighted averages* (OWA) are defined by

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

where $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_p)$ is a weighting vector such that $\sum_{i=1}^{p} \lambda_i = 1$ and $\lambda_i \geq 0\ \forall i$. The $\boldsymbol{\lambda}$ is a parameter of the model providing a control on the relative importance attached to bad and good performances.

– The *Choquet integrals* are defined by

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

where v is a set function, named *capacity*, defined on 2^C such that $v(\emptyset) = 0$, $v(C) = 1$ and $v(X) \leq v(Y)$ for all $X \subseteq Y \subseteq C$. The capacity v models the importance attached to any coalition of criteria and makes it possible to represent positive or negative interactions between criteria [6].

The parameters of the aggregation functions introduced above (i.e., the weighting vector or the capacity function) must be tuned to the DM's value system in order to make relevant recommendations. Our aim is to collect preference statements to progressively reduce the uncertainty about these parameters to make a more cautious recommendation.

In order to assess parameters $\boldsymbol{\lambda}$ or v with Bayesian linear regression methods, it is important to remark that OWA and Choquet admit linear reformulations of the form $f_{\mathbf{w}}(\mathbf{a}) = \sum_{i=1}^{q} w_i g_i(\mathbf{a})$, where $\mathbf{w} = (w_1, \ldots, w_q)$ is a weighting vector and $\{g_1(\mathbf{a}), \ldots, g_q(\mathbf{a})\}$ is a generating set of non-linear functions defined from criterion values.

[1] Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, LIP6, F-75005 Paris, France, email: name.surname@lip6.fr

For OWA with decreasing weights, we can use p cumulative functions defined by $g_i(\mathbf{a}) = \sum_{k=1}^{i} a_{(k)}$, and positive weights $w_i = \lambda_i - \lambda_{i+1}$ for $i < p$ and $w_p = \lambda_p$. It can also be shown that another linear reformulation exists for 2-additive Choquet integrals with $\mathbf{g}_i(\mathbf{a}) = C_{v_i}(\mathbf{a})$, using specific 0-1 capacities v_i (unanimity games and some of their conjugates). Details are omitted due to the space constraint.

3 Elicitation by expected regret minimization

We propose a Bayesian approach where a prior Gaussian distribution 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. Let us now introduce the notion of expected regret more formally, for an aggregation function $f_{\mathbf{w}}$.

Definition 1. *Given a probability distribution 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:*

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

Definition 2. *Given a set $\mathcal{A}$ of alternatives and a probability distribution p on the possible values of $\mathbf{w}$, the max expected regret of $\mathbf{a} \in \mathcal{A}$ is defined by* $\mathrm{MER}(\mathbf{a}, \mathcal{A}, p) = \max_{\mathbf{b} \in \mathcal{A}} \mathrm{PER}(\mathbf{a}, \mathbf{b}, p)$

Definition 3. *Given a set $\mathcal{A}$ of alternatives and a probability distribution p on the possible values of $\mathbf{w}$, the minimax expected regret is defined by* $\mathrm{MMER}(\mathcal{A}, p) = \min_{\mathbf{a} \in \mathcal{A}} \mathrm{MER}(\mathbf{a}, \mathcal{A}, p)$.

The incremental decision procedure is described in Algorithm 1, where $\mathbf{a}^{(\mathbf{i})} \in \mathcal{A}$ denotes an alternative achieving the $\mathrm{MMER}(\mathcal{A}, p)$ value at the current iteration i and $\mathbf{b}^{(\mathbf{i})}$ is its stronger opponent.

Algorithm 1: Incremental Decision Making

Input: $\mathcal{A}$: set of alternatives, δ: acceptance threshold;
p^0: prior distribution on vectors $\mathbf{w}$.
Output: $\mathbf{a}^*$: best recommendation in $\mathcal{A}$

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

The probability distribution $p(\mathbf{w})$ is updated in Line 7 using the latent variable approach introduced by [1] for binary *probit* regression, where the noise is assumed to be Gaussian. The latent variable corresponds to a preference intensity (utility gap) that is not observable but that makes it possible to explain the DM's answer $y^{(i)}$.

4 Numerical tests

The results are averaged over 20 randomly generated instances with 100 Pareto optimal alternatives and 5 criteria. The DM's answers are simulated using an hidden OWA model perturbated by a Gaussian noise $\mathcal{N}(0, \sigma)$. Hence, the percentage of errors made by the DM can be controlled by varying σ. The experiments show that the algorithm converges in a dozen of queries. The curves in Figure 1 show, at any step, the improvement of the quality of the current MMER-optimal alternative, for various values of σ. This quality is measured by the rank of the alternative according to the hidden DM preference model.

We remark that the average rank of the recommended alternative is never beyond 5 (over 100 alternatives), thus the procedure exhibits a good tolerance to errors. Concerning the computation times between two queries, it takes about 5 seconds for OWA with decreasing weights, and about 35 seconds for 2-additive Choquet integrals.

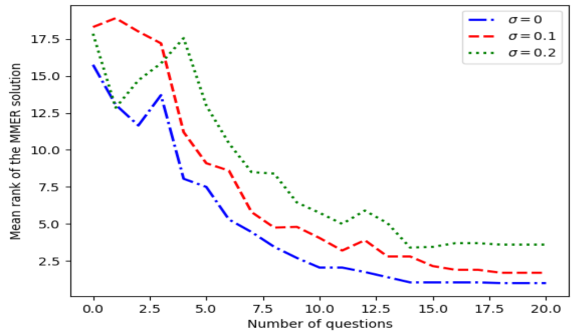

Figure 1. Mean rank of the recommendation (OWA).

REFERENCES

[1] James H Albert and Siddhartha Chib, 'Bayesian analysis of binary and polychotomous response data', *Journal of the American statistical Association*, **88**(422), 669–679, (1993).

[2] 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).

[3] 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).

[4] Urszula Chajewska, Daphne Koller, and Ronald Parr, 'Making rational decisions using adaptive utility elicitation.', in *Proceedings of AAAI-00*, pp. 363–369, (2000).

[5] Wei Chu and Zoubin Ghahramani, 'Preference learning with gaussian processes', in *Proceedings of ICML 2005*, pp. 137–144, (2005).

[6] Michel Grabisch and Christophe Labreuche, 'A decade of application of the Choquet and Sugeno integrals in multi-criteria decision aid', *Annals of Operations Research*, **175**(1), 247–286, (2010).

[7] Shengbo Guo and Scott Sanner, 'Multiattribute bayesian preference elicitation with pairwise comparison queries', in *Proceedings of NIPS*, pp. 396–403, (2010).

[8] Vu Ha and Peter Haddawy, 'Problem-focused incremental elicitation of multi-attribute utility models', in *Proceedings of UAI-97*, pp. 215–222. Morgan Kaufmann Publishers Inc., (1997).

[9] Ahti A Salo and Raimo P Hämäläinen, 'Preference programming through approximate ratio comparisons', *European Journal of Operational Research*, **82**(3), 458–475, (1995).

[10] 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).

[11] Paolo Viappiani and Craig Boutilier, 'Optimal bayesian recommendation sets and myopically optimal choice query sets', in *Proceedings of NIPS*, pp. 2352–2360, (2010).

[12] Tianhan Wang and Craig Boutilier, 'Incremental utility elicitation with the minimax regret decision criterion', in *IJCAI-03*, pp. 309–318, (2003).

[13] Chelsea C White, Andrew P Sage, and Shigeru Dozono, 'A model of multiattribute decision making and trade-off weight determination under uncertainty', *IEEE Transactions on Systems, Man, and Cybernetics*, **14**(2), 223–229, (1984).

[14] Ronald R Yager, 'On ordered weighted averaging aggregation operators in multicriteria decisionmaking', *IEEE Transactions on systems, Man, and Cybernetics*, **18**(1), 183–190, (1988).