



HAL
open science

A metaheuristic for inferring a ranking model based on multiple reference profiles

Arwa Khannoussi, Alexandru-Liviu Olteanu, Patrick Meyer, Bastien Padeloup

► **To cite this version:**

Arwa Khannoussi, Alexandru-Liviu Olteanu, Patrick Meyer, Bastien Padeloup. A metaheuristic for inferring a ranking model based on multiple reference profiles. *Annals of Mathematics and Artificial Intelligence*, 2024, 10.1007/s10472-024-09926-w . hal-04017642v2

HAL Id: hal-04017642

<https://imt-atlantique.hal.science/hal-04017642v2>

Submitted on 9 Feb 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.

A metaheuristic for inferring a ranking model based on multiple reference profiles

Arwa Khannoussi¹, Alexandru-Liviu Olteanu³, Patrick Meyer^{2*} and Bastien Padeloup²

¹IMT Atlantique, LS2N, UMR CNRS 6004, F-4430 Nantes, France.

²IMT Atlantique, Lab-STICC, UMR CNRS 6285, F-29238 Brest, France.

³Lab-STICC, UMR 6285, CNRS, Université Bretagne Sud, Lorient, France.

*Corresponding author(s). E-mail(s):

patrick.meyer@imt-atlantique.fr;

Contributing authors: arwa.khannoussi@imt-atlantique.fr;

alexandru.olteanu@univ-ubs.fr;

bastien.padeloup@imt-atlantique.fr;

Abstract

In the context of Multiple Criteria Decision Aiding, decision makers often face problems with multiple conflicting criteria that justify the use of preference models to help advancing towards a decision. In order to determine the parameters of these preference models, preference elicitation makes use of preference learning algorithms, usually taking as input holistic judgments, i.e., overall preferences on some of the alternatives, expressed by the decision maker. Tools to achieve this goal in the context of a ranking model based on multiple reference profiles are usually based on mixed-integer linear programming, Boolean satisfiability formulation or metaheuristics. However, they are usually unable to handle realistic problems involving many criteria and a large amount of input information. We propose here an evolutionary metaheuristic in order to address this issue. Extensive experiments illustrate its ability to handle problem instances that previous proposals cannot.

Keywords: multi-criteria decision aiding, preference elicitation and learning, ranking problem, reference profiles, meta-heuristic

1 Introduction

Difficult decisions usually involve multiple, often conflicting, objectives or criteria over a set of alternatives. In *Multiple Criteria Decision Aiding* (MCDA), we distinguish three types of decision problems [20]: 1. *Choice* refers to selecting the best alternative or the best set of alternatives; 2. *Ranking* aims to order all alternatives from the best one to the worst; 3. *Sorting* seeks to assign alternatives to one or several predefined preferentially ordered categories. In this work, we focus on the ranking problem, and, more specifically, the *Ranking based on Multiple reference Profiles* (RMP) [18, 6] preference model. In this model, pairs of alternatives are not compared directly, but through a set of underlying reference profiles. This process allows to construct a preference relation on the set of alternatives. The RMP model belongs to the class of MCDA methods that are based on outranking relations [19]. In this work, we restrict ourselves to a specific case of the RMP model, called *Simple RMP* (SRMP) [18], in which the importance of criteria is represented by additive weights.

The use of reference profiles is widely spread among MCDA methods, and is based on psychological evidence [21] which suggests that *decision makers* (DMs) often base their decisions on so-called references which correspond to their current expectations on the decision problem. For ranking problems we can mention the TOPSIS [10] method, which evaluates alternatives using distances to the ideal and anti-ideal points, to be minimized and maximized respectively. The MACBETH [2] method also uses two reference levels on each criterion, corresponding to “good” and “neutral” evaluations, in order to elicit a value based model. For sorting problems, we can mention the ELECTRE TRI [19] method and its multiple variants [15, 1, 4, 5] which compare alternatives to limiting or central category profiles.

In order to be used in practice, the parameters of the SRMP ranking model need to be determined so that the model accurately reflects the perspective of the DM. The interaction between an analyst and a DM in order to set the parameters of a preference model is called the *preference elicitation* process. The direct elicitation approach [22] requires direct involvement of the DM and should use well-designed questions formulated in terms of the primitives of the underlying model. The direct rating approach requires the DM to give numerical values for the model parameters. Finally the indirect approach uses holistic information given by the DM, i.e. preferences on some of the alternatives, in order to learn the model parameters [11] through a preference learning algorithm. The direct elicitation and direct rating approaches are usually difficult to apply in practice, as the DM has to perfectly understand how the preference model works [7] (model-driven approach), while the indirect approach does not require this expertise and simply asks the DM to express judgments on the desired output of the method for a few alternatives only (data-driven approach). In the case of the SRMP model, this holistic information corresponds to pairwise comparisons of alternatives.

Previous work on inferring the parameters of an SRMP model [17] involves an exact resolution approach, more precisely a mixed integer linear optimization problem (MILP), that requires significant computational resources and time, even for a small number of input pairwise comparisons. Belahcène *et al.* [3] propose to use a Boolean satisfiability approach (SAT) to learn the more general RMP model. They reduce execution time compared to [17] and allow for larger sets of input pairwise comparisons to be used. A metaheuristic algorithm, integrating linear programming and a heuristic in its design, is proposed by Liu *et al.* [16] to infer the parameters of an SRMP model. This approach is faster than the previous ones, allows to handle larger sets of pairwise comparisons of alternatives and more complex SRMP models (in terms of number of criteria and profiles), but this leads to a loss in the accuracy on the learning dataset. The topic of incrementally learning the parameters of the SRMP model through repeated interactions with the DM has also been recently addressed by Khannoussi *et al.* [14, 12, 13].

We extend these works by proposing an evolutionary metaheuristic, more specifically a genetic algorithm for learning the parameters of an SRMP model. We show that this proposal allows to deal with large instances in seconds for problems that were considered as difficult by the previous approaches. This allows the use of this learning algorithm in real-life problems. Despite not always reaching 100% accuracy on the input data, the inferred models have a high generalization ability on unseen data, similar to those reported in the previous articles, but in far less time.

The paper is structured in the following way: Section 2 introduces the reader to the SRMP method. Section 3 presents the proposed preference learning algorithm, and Section 4 the tuning of its hyperparameters. Section 5 provides a numerical analysis of its performance on problem instances of different sizes and input information of varying quality. Finally, Section 6 concludes with several final remarks and perspectives for future work.

2 SRMP: a Ranking model based on Multiple Profiles

We consider a finite set of A distinct alternatives, and a fixed set of C criteria. We note $\mathbf{A} \in [0, 1]^{A \times C}$ the matrix of evaluations of all alternatives on all criteria, *i.e.*, for an alternative a and a criterion c , $\mathbf{A}[a, c] \in [0, 1]$. We use the “:” operator to span the entire dimension of an associated matrix. Hence $\mathbf{A}[:, c]$ is the c^{th} column of \mathbf{A} , and $\mathbf{A}[a, :]$ its a^{th} row.

Without loss of generality, we restrict all evaluations to the interval $[0, 1]$, and assume that higher evaluations of alternatives on criteria are preferred to lower ones.

2.1 Parameters of the SRMP model

The SRMP preference model [18] is characterized by the following preference parameters:

4 *A metaheuristic for inferring an SRMP model*

- P reference profiles, each evaluated on all criteria, denoted by a matrix $\mathbf{P} \in [0, 1]^{P \times C}$. Profiles satisfy the dominance rule defined by:

$$\forall c \in \{1, \dots, C\}, \forall p \in \{1, \dots, P-1\} : \mathbf{P}[p, c] \leq \mathbf{P}[p+1, c]; \quad (1)$$

- A lexicographic order on the profiles, denoted by a vector of integers \mathbf{o} which represents the permutation of the indexes of the profiles. This order represents the importance of the reference points and defines thus the sequence in which each alternative is compared to the profiles.
- A vector of weights $\mathbf{w} \in]0, 1[^C$ associated with the criteria, satisfying the following constraint:

$$\sum_c \mathbf{w}[c] = 1. \quad (2)$$

An SRMP model M is therefore described by the tuple $(\mathbf{P}, \mathbf{o}, \mathbf{w})$.

Figure 1 represents an example of an SRMP model, for $C = 4$ criteria and $P = 3$ reference profiles. Each criterion is represented by a vertical axis, on which the preferred performances are plotted towards the top (in the direction of the arrows). The dotted lines are the profiles, whose lexicographic order (\mathbf{o}) is shown on the left. The weights (\mathbf{w}) are shown below each axis.

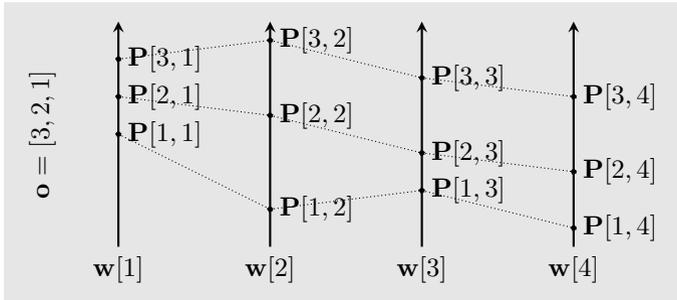


Fig. 1: Representation of an SRMP model for $C = 4$ criteria and $P = 3$ reference profiles.

2.2 Comparing alternatives, given an SRMP model

An SRMP model can be used to provide a ranking on the set of alternatives by comparing them sequentially, in pairwise manner, to the reference profiles considered in the order given by \mathbf{o} .

More precisely, two alternatives a_1 and a_2 are in a strict preference relation w.r.t. profile p , denoted $a_1 \succ_p a_2$, iff:

$$\sum_{\mathbf{A}[a_1, c] \geq \mathbf{P}[p, c]} \mathbf{w}[c] > \sum_{\mathbf{A}[a_2, c] \geq \mathbf{P}[p, c]} \mathbf{w}[c]. \quad (3)$$

Alternatively, they are in an indifference relation w.r.t. profile p , denoted $a_1 \sim_p a_2$, iff:

$$\sum_{\mathbf{A}[a_1, c] \geq^c \mathbf{P}[p, c]} \mathbf{w}[c] = \sum_{\mathbf{A}[a_2, c] \geq^c \mathbf{P}[p, c]} \mathbf{w}[c]. \quad (4)$$

An overall relation between two alternatives a_1 and a_2 is constructed by sequentially considering the profiles using the lexicographic order \mathbf{o} .

A strict preference relation $a_1 \succ a_2$ occurs when $a_1 \succ_{\mathbf{o}[1]} a_2$, or when $a_1 \sim_{\mathbf{o}[1]} a_2$ and $a_1 \succ_{\mathbf{o}[2]} a_2$, or \dots or $a_1 \sim_{\mathbf{o}[p]} a_2, \forall p \in \{1, \dots, P-1\}$ and $a_1 \succ_{\mathbf{o}[P]} a_2$. An indifference relation $a_1 \sim a_2$ occurs when $a_1 \sim_p a_2, \forall p \in \{1, \dots, P\}$.

It has been shown that \succ and \sim define a preorder on the set of alternatives [18].

3 The preference inference algorithm

The goal of the preference inference algorithm that we propose is to determine the parameters of the SRMP model, i.e., the reference profiles, the lexicographic order and the weights, given the holistic preferences expressed by a DM on pairs of alternatives.

The proposed algorithm uses the classical ingredients of a genetic algorithm [9]. A genetic algorithm is a metaheuristic from the class of evolutionary algorithms, which aims at generating good solutions to optimization problems. It is inspired by Charles Darwin's theory of natural evolution, and mimics the natural selection process, where the "fittest" individuals (the solutions) of a population are selected for reproduction in order to produce offsprings of the next generation.

Algorithm 1 sums up the elements introduced more precisely in the following sections. Implementation details and codes to reproduce our experiments are publicly available at https://github.com/BastienPasdeloup/learn_srmp.

3.1 Solution

In our context, a *solution* in the genetic algorithm consists of the parameters of an SRMP model M , i.e., the reference profiles \mathbf{P} , criteria weights \mathbf{w} and lexicographic order \mathbf{o} , as described in Section 2.

It is worth noting that an SRMP model requires the knowledge of the value of P , that determines the number of reference profiles in a solution. This value is unknown in practice and has thus to be given as an input argument of Algorithm 1.

3.2 Fitness

Let $\mathcal{D}_{\text{TRAIN}}$ be a set of pairwise comparisons of alternatives provided by the DM, i.e., each element of $\mathcal{D}_{\text{TRAIN}}$ is a tuple (a_1, a_2, \odot) , where a_1 and a_2 are two distinct alternatives, and $\odot \in \{\succ, \prec, \sim\}$ is the preference relation $a_1 \odot a_2$ provided by the DM.

6 *A metaheuristic for inferring an SRMP model*

For a given solution M in the genetic algorithm, the number of pairs of alternatives that M is able to compare similarly to the DM therefore gives a quality measure of M . The higher this value, the more representative of the DM is M . As a consequence, in our genetic algorithm, we use the following *fitness function* to evaluate a solution:

$$\text{fitness}(M, \mathcal{D}_{\text{TRAIN}}) = \frac{1}{|\mathcal{D}_{\text{TRAIN}}|} \sum_{\{(a_1, a_2, \odot)\} \in \mathcal{D}_{\text{TRAIN}}}} \begin{cases} 1 & \text{if } \odot_M = \odot \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

where $\odot_M \in \{\succ, \prec, \sim\}$ gives the ordering $a_1 \odot_M a_2$ obtained thanks to M .

Algorithm 1 GA_SRMP (\mathcal{T}, S, B, R, P)

```

1: Input
2:  $\triangleright \mathcal{T}$ : Training set used to learn the SRMP model
3: crossover( $\cdot, \cdot$ ): Crossover operator (applied on two solutions), returns a
   child solution
4: mutation( $\cdot$ ): Mutation operator (applied on one solution), returns a
   mutated solution
5:  $\triangleright S$ : Population size of the genetic algorithm
6:  $\triangleright B$ : Ratio of best solutions in the population kept for next iteration
7:  $\triangleright R$ : Ratio of random solutions injected at next iteration
8:  $\triangleright P$ : Number of reference profiles to create in solution
9: Output
10:  $\triangleright \widehat{M}$ : Best solution found

```

```

11:  $pop_{it} \leftarrow \text{initialize\_pop}(S, P)$   $\triangleright$  Initial population of random solutions
12: while not stopping\_condition( $pop_{it}$ ) do  $\triangleright$  Check stopping condition
13:    $pop_{it+1} \leftarrow \text{prepare\_pop}(\mathcal{T}, pop_{it}, S, B, R)$   $\triangleright$  Initialize next pop
14:   while  $\text{size}(pop_{it+1}) \neq S$  do  $\triangleright$  Fill with generated solutions
15:      $M_1, M_2 \leftarrow \text{select}_2(pop_{it})$   $\triangleright$  Select 2 parents using roulette wheel
16:      $M \leftarrow \text{crossover}(M_1, M_2)$   $\triangleright$  Apply crossover
17:      $M \leftarrow \text{mutation}(M)$   $\triangleright$  Apply mutation
18:     if  $\text{fitness}(M, \mathcal{T}) > \min(\text{fitness}(M_1, \mathcal{T}), \text{fitness}(M_2, \mathcal{T}))$  then
19:        $pop_{it+1} \leftarrow pop_{it+1} \cup \{M\}$   $\triangleright$  Keep if better than 1 parent
20:      $pop_{it} \leftarrow pop_{it+1}$   $\triangleright$  Go to next iteration
return  $\text{solution\_with\_best\_fitness}(pop_{it})$   $\triangleright$  Return  $\widehat{M}$ 

```

3.3 Initialization

The genetic algorithm *initializes* a population of solutions and then improves it through repetitive applications of *selection*, *mutation* and *crossover* operators. The *initial population* is filled by randomly generating valid solutions, *i.e.*, that respect Equations 1 and 2, as well as definition domains introduced in Section 2.

3.4 Main loop

The population at a given iteration is constructed with some of the best solutions in the population at the end of the previous iteration, in addition to some new random valid solutions in order to maintain diversity (`prepare_pop`). The rest of the population is obtained through a reproduction method, where some *parent solutions* are *selected* to create *child solutions*. Individual solutions are selected based on the value of their fitness using a *roulette wheel* procedure, which tends to prefer “good” solutions in the selection process, while allowing for “less good” solutions to be selected sometimes with a probability proportional to their fitness.

To generate new solutions, crossover and mutation operators are used. A *crossover* operator produces child solutions from the selected parent solutions, by keeping some characteristics of the parents. The children are then transformed through a *mutation* operator to diversify the population, and are finally inserted in the population if they improve the fitness of at least one of the parents.

After each operation – crossover or mutation –, solutions are repaired to ensure their validity. In more detail, the following corrections are applied:

1. For each criterion c , entries of $\mathbf{P}[:, c]$ are sorted. Entries of \mathbf{P} are then clipped to belong to $[0, 1]$, *i.e.*, values outside the interval are updated to the closest interval edge;
2. Entries of \mathbf{w} are first clipped to belong to $]0, 1[$. Then, for each criterion c , $\mathbf{w}[c] \leftarrow \frac{\mathbf{w}[c]}{\|\mathbf{w}\|_1}$, where $\|\cdot\|_1$ denotes the ℓ_1 norm.

3.5 Proposed crossover operators

Some key elements of a genetic algorithm are the operators it uses. This section describes the crossover operators introduced in this work. These operators can all be used in Algorithm 1 as `crossover(·, ·)`. They are illustrated in Figure 2.

Let $M_1 = (\mathbf{P}_1, \mathbf{o}_1, \mathbf{w}_1)$ and $M_2 = (\mathbf{P}_2, \mathbf{o}_2, \mathbf{w}_2)$ be two selected parent solutions with P reference profiles, and let M be a child solution as generated by the presented crossover operators. In practice, the implementation also generates a second child complementary to M . We do not detail more for simplicity of exposition, and refer the interested reader to the provided Github repository:

- `swap_profiles`: Exchanges profiles between two solutions:

$$M \leftarrow (\mathbf{P}_2, \mathbf{o}_1, \mathbf{w}_1);$$

- `swap_orders`: Exchanges the lexicographic orders between two solutions:

$$M \leftarrow (\mathbf{P}_1, \mathbf{o}_2, \mathbf{w}_1);$$

- `swap_weights`: Exchanges weights between two solutions:

$$M \leftarrow (\mathbf{P}_1, \mathbf{o}_1, \mathbf{w}_2);$$

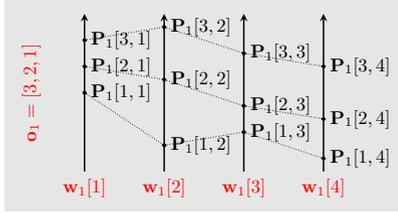
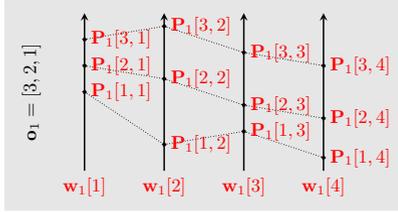
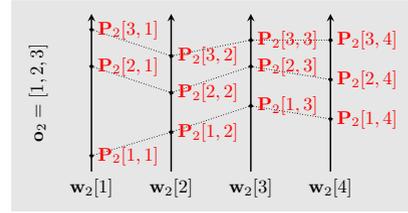
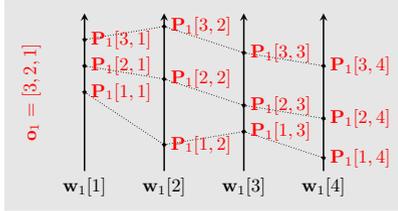
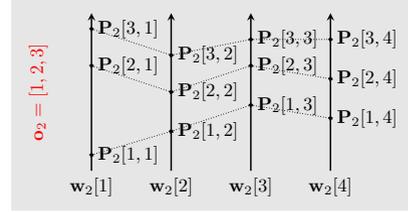
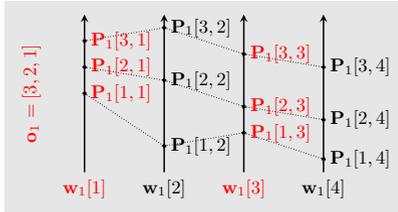
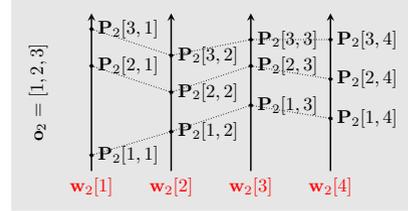
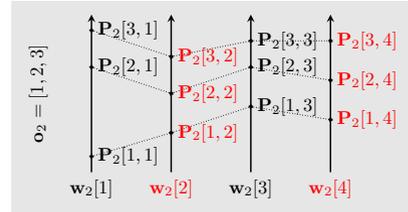
8 *A metaheuristic for inferring an SRMP model*(a) `swap_profiles`(b) `swap_orders`(c) `swap_weights`(d) `mix_criteria_and_weights`

Fig. 2: Proposed crossover operators. Elements in red are assembled to build a child. Note that elements in black can also be assembled to produce a second child, as mentioned in Section 3 and done in our implementation.

- `mix_criteria_and_weights`: Exchanges the profile values and associated weights on a random subset of criteria between two solutions. Let $\mathcal{S} \subset \{1, \dots, C\}$. Let $\mathbf{P} \in \mathbb{R}^{P \times C}$ such that $\forall p, \forall c : \mathbf{P}[p, c] = \mathbf{P}_1[p, c]$ if $c \in \mathcal{S}$, and $\mathbf{P}[p, c] = \mathbf{P}_2[p, c]$ otherwise. Additionally, let $\mathbf{w} \in \mathbb{R}^C$ such that

$\forall p, \forall c : \mathbf{w}[c] = \mathbf{w}_1[c]$ if $c \in \mathcal{S}$, and $\mathbf{w}[c] = \mathbf{w}_2[c]$ otherwise:

$$M \leftarrow (\mathbf{P}, \mathbf{o}_1, \mathbf{w}) .$$

All of these operators copy entire parts (profiles, order, weights) of both parent solutions into the child solutions, with the aim, therefore, of potentially transmitting good characteristics into future generations.

3.6 Proposed mutation operators

Once created using a crossover operator, child solutions can be altered using a mutation. This section describes the mutation operators introduced in this work. These operators are provided to Algorithm 1 as `mutation(·)`. They are illustrated in Figure 3.

Let $M = (\mathbf{P}, \mathbf{o}, \mathbf{w})$ be the solution to be mutated:

- **random_profile_perturbation:** Applies an independent and identically distributed (i.i.d.) random perturbation of amplitude $\lambda_{\text{RPP}} \in]0, 1[$ to all the profile values. Let $\mathbf{X} \in \mathbb{R}^{P \times C}$ such that $\forall p, \forall c : \mathbf{X}[p, c] \sim \mathcal{U}(-1, 1)$:

$$\forall p, \forall c : \mathbf{P}[p, c] \leftarrow \mathbf{P}[p, c] + \lambda_{\text{RPP}} \mathbf{X}[p, c] ;$$

- **random_weights_perturbation:** Applies a (i.i.d.) random perturbation of amplitude $\lambda_{\text{RWP}} \in]0, 1[$ to the weights of the criteria. Let $\mathbf{x} \in \mathbb{R}^C$ such that $\forall c : \mathbf{x}[c] \sim \mathcal{U}(-1, 1)$:

$$\forall c : \mathbf{w}[c] \leftarrow \mathbf{w}[c] + \lambda_{\text{RWP}} \mathbf{x}[c] ;$$

- **shrink_profiles:** Shrinks the space between profiles by a factor $\lambda_{\text{SP}} \in]0, 1[$ to bring them closer to one another, while preserving the original mean of all profiles per criterion. Let $\mathbf{m}_1 \in \mathbb{R}^C$ such that $\forall c : \mathbf{m}_1[c] = \frac{1}{P} \|\mathbf{P}[:, c]\|_1$. We proceed in two steps:

$$\forall p, \forall c : \mathbf{P}[p, c] \leftarrow \frac{1}{1 + \lambda_{\text{SP}}} \mathbf{P}[p, c] .$$

Then, let $\mathbf{m}_2 \in \mathbb{R}^C$ such that $\forall c : \mathbf{m}_2[c] = \frac{1}{P} \|\mathbf{P}[:, c]\|_1$ computed after this operation. We then update profile values as follows:

$$\forall p, \forall c : \mathbf{P}[p, c] \leftarrow \mathbf{P}[p, c] - \mathbf{m}_2[c] + \mathbf{m}_1[c] ;$$

- **expand_profiles:** Conversely, expands the space between profiles by a factor $\lambda_{\text{EP}} \in]0, 1[$ to push them apart, while preserving the original mean of all profiles per criterion. Let $\mathbf{m}_1 \in \mathbb{R}^C$ such that $\forall c : \mathbf{m}_1[c] = \frac{1}{P} \|\mathbf{P}[:, c]\|_1$. We proceed in two steps:

$$\forall p, \forall c : \mathbf{P}[p, c] \leftarrow (1 + \lambda_{\text{EP}}) \mathbf{P}[p, c] .$$

Then, let $\mathbf{m}_2 \in \mathbb{R}^C$ such that $\forall c : \mathbf{m}_2[c] = \frac{1}{P} \|\mathbf{P}[:, c]\|_1$ computed after this operation. We then update profile values as follows:

$$\forall p, \forall c : \mathbf{P}[p, c] \leftarrow \mathbf{P}[p, c] - \mathbf{m}_2[c] + \mathbf{m}_1[c] ;$$

- `partially_reverse_order`: Reverses a random part of the lexicographic order. Let i and j be random integers in $\{1, \dots, P\}$ with $i < j$:

$$\forall k \in [i, j] : \mathbf{o}[k] \leftarrow \mathbf{o}[j - k + i] .$$

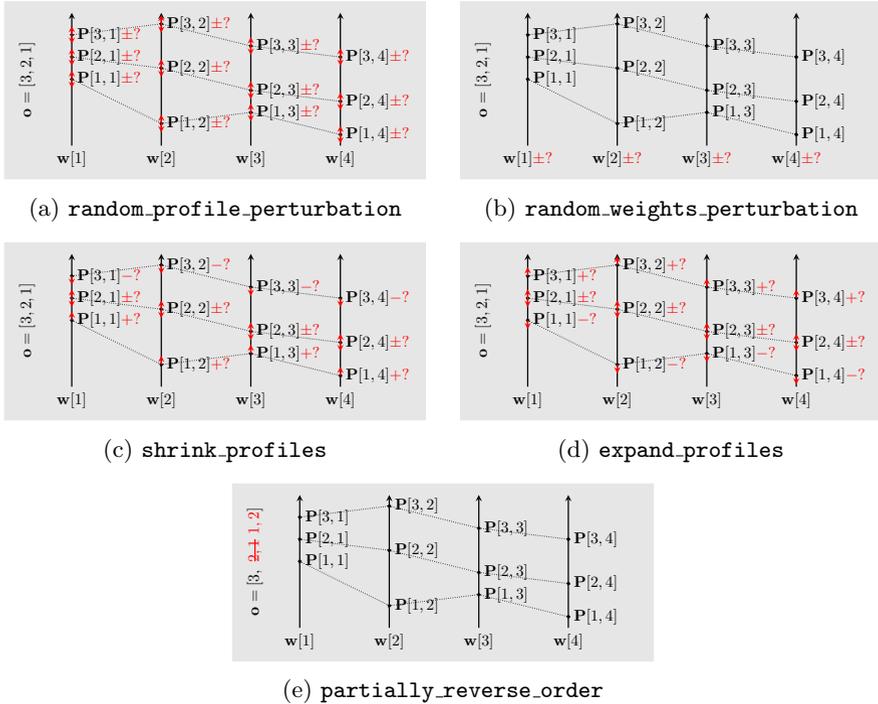


Fig. 3: Proposed mutation operators. Elements in red highlight the alterations operated on the child. Symbols $?$ indicate a random evolution of the parameter, with magnitude controlled by the corresponding perturbation parameter (which we write λ . in the sequel).

3.7 Stopping criterion

Multiple stopping criteria can be considered depending on the target application. For example, the algorithm could be stopped after a certain number of iterations, a maximum execution time, when the value of the fitness function

for the best solution in the current population reaches a given value, or when the fitness function of the best solution does not change for a certain number of iterations or a given interval of time.

In order to favor solutions that maximize accuracy, we choose in this work to use a stopping criterion which checks that no evolution of the K best (*i.e.*, with highest fitness) solutions for a given number of consecutive iterations occurs, independently of execution time.

In all of our experiments below the value of K is chosen as $K = 0.1BS$, corresponding to no changes among the top 10% within the BS best solutions in the population, for 50 consecutive iterations. These values were established empirically by trial and error on a large number of data sets.

4 Preliminary experiments

4.1 Considered settings

A first step before performing experiments to evaluate the performance of the proposed genetic algorithm consists in determining the values of its numerous hyperparameters. Additionally, we are interested in evaluating which of the crossover and mutation operators introduced in Section 3 are indeed useful, in the sense that using them helps reaching models with high accuracy.

In the experiments described in this section, we have chosen to focus on problems with the following characteristics:

- Number of reference profiles (ground truth value of P): 3;
- Number of criteria (C): 11;
- Number of alternatives (A): 50;
- Size of the training set ($|\mathcal{D}_{\text{TRAIN}}|$): 100.

Decision problems encountered in practice often have this type of characteristics, which leads us to carry out the study of the hyperparameters on these data sizes.

The evaluations of the A alternatives on the C criteria are randomly drawn from a uniform distribution between 0 and 1, and stored in matrix $\mathbf{A} \in [0, 1]^{A \times C}$.

We consider in this section 10 distinct random ground truth DMs, that we name $\{M_1^*, \dots, M_{10}^*\}$. The models associated with each of these DMs are generated as follows:

- all evaluations in \mathbf{P} are randomly drawn from a uniform distribution between 0 and 1 and then sorted in increasing order column-wise;
- the vector of weights \mathbf{w} is constructed using the approach from [8];
- the lexicographic order \mathbf{o} is randomly generated.

The training set is constructed by randomly selecting pairs of alternatives from the set of A alternatives and generating the relation between them using a given DM model. We denote $\mathcal{D}_{\text{TRAIN}}^i$ the training set using the i^{th} ground truth model M_i^* .

Also, we introduce a new matrix $\mathbf{A}' \in [0, 1]^{A' \times C}$ of $A' = 300$ unseen alternatives, *i.e.*, distinct from those in \mathbf{A} . Similarly to $\mathcal{D}_{\text{TRAIN}}^i$, we note $\mathcal{D}_{\text{TEST}}^i$

the set of all possible comparisons of distinct alternatives in \mathbf{A}' , compared using the associated ground truth DM M_i^* .

4.2 Grid search

For each pair of a mutation and a crossover operators, we perform a grid search – *i.e.*, exhaustive combination – of the following values for the corresponding operators:

- Population size (S in Algorithm 1): {100, 150, 200, 250, 300};
- Ratio of best solutions transmitted to next iteration (B in Algorithm 1): {0.1, 0.2, 0.3, 0.4, 0.5};
- Ratio of new random solutions per iteration (R in Algorithm 1): {0.1, 0.2, 0.3, 0.4, 0.5};
- Number of reference profiles to learn (P in Algorithm 1): 3.
- λ_{RPP} (if applicable): {0.1, 0.2, 0.3, 0.4, 0.5};
- λ_{RWP} (if applicable): {0.1, 0.2, 0.3, 0.4, 0.5};
- λ_{SP} (if applicable): {0.2, 0.4, 0.6, 0.8, 1};
- λ_{EP} (if applicable): {0.2, 0.4, 0.6, 0.8, 1}.

Note that P in Algorithm 1 is chosen to match the actual number of reference profiles of the ground truth DM to learn. In practical cases, we want to model real-life DMs, and no such *true* P exists. Therefore, a choice of P is needed. This point is discussed in more details in Section 5.

These experiments are repeated for each pair of a crossover and a mutation operator, in isolation from all other operators. Each configuration is evaluated on all the ground truth DMs to provide average results. Therefore, for each pair of a mutation and a crossover operator, we have run $5*5*5*5*10 = 6250$ experiments (except when the mutation operator is `partially_reverse_order` for which we have run $5*5*5*10 = 1250$ experiments, since there is no associated λ). All together, we have performed $16*6250 + 4*1250 = 105000$ experiments, corresponding to 10500 different configurations. Let $\Omega = \{\omega_1, \dots, \omega_{10500}\}$ be the set of all these configurations.

For each configuration $\omega \in \Omega$, we therefore obtain 10 estimated DMs $\{\widehat{M}_1^\omega, \dots, \widehat{M}_{10}^\omega\}$ of the ground truth DMs $\{M_1^*, \dots, M_{10}^*\}$. We measure the following quantities:

- The *train accuracy* of \widehat{M}_i^ω , given using Equation 5 by:

$$\text{train_accuracy}(\widehat{M}_i^\omega) = \text{fitness}(\widehat{M}_i^\omega, \mathcal{D}_{\text{TRAIN}}^i).$$

From this, we derive the *average train accuracy* of ω :

$$\text{average_train_accuracy}(\omega) = \frac{1}{10} \sum_{i=1}^{10} \text{train_accuracy}(\widehat{M}_i^\omega); \quad (6)$$

- The *test accuracy* of \widehat{M}_i^ω , given using Equation 5 by:

$$\text{test_accuracy}(\widehat{M}_i^\omega) = \text{fitness}(\widehat{M}_i^\omega, \mathcal{D}_{\text{TEST}}^i).$$

Similarly, we derive the *average test accuracy* of ω :

$$\text{average_test_accuracy}(\omega) = \frac{1}{10} \sum_{i=1}^{10} \text{test_accuracy}(\widehat{M}_i^\omega). \quad (7)$$

4.3 Hyperparameters tuning

Given the outputs of the experiments described above, we select adequate values for the hyperparameters, and pertinent pairs of crossover and mutation operators to use. To do so, we proceed as follows:

1. Sort configurations in Ω by decreasing average train accuracy;
2. Initialize another set Ω^* with the first configuration from Ω . Elements in Ω^* are considered to be significantly better than the remaining configurations from Ω ;
3. For any remaining configuration $\omega \in \Omega$:
 - 3.1 Compare ω to all configurations $\omega^* \in \Omega^*$ using a paired samples t-test (with a standard significance level of 0.05) on the associated sets $\{M_1^\omega, \dots, M_{10}^\omega\}$ and $\{M_1^{\omega^*}, \dots, M_{10}^{\omega^*}\}$;
 - 3.2 Add ω to Ω^* if it is not worse (according to the test) than any configuration in Ω^* ;
 - 3.3 Stop if Ω^* does not change for a given number of iterations (20 in this case), otherwise continue.

In our experiment, after this procedure, we were left with $|\Omega^*| = 18$ configurations that cannot be considered significantly better than each other. For these solutions, we obtain an average train accuracy of 0.991 and an average test accuracy of 0.975. Regarding execution time, we observe an average of 32.44 seconds per experiment on a machine with Intel Xeon Gold CPU @ 2.7GHz \times 56, with 256GB RAM, with the provided Python 3.6 implementation of Algorithm 1.

Among solutions in Ω^* , we count the number of occurrences of each attributions of a value to a parameter in the grid search. Results are presented in Table 1.

Table 1 reveals some clear conclusions on the hyperparameters:

- *S*: The higher the better. We retain a value of $S = 300$ for subsequent experiments. Higher values may be used, but this has a consequence on execution time, and train accuracy is already close to 1 in all simulations;
- *B*: Again, high values are preferred. We retain a value of $B = 0.4$;
- *R*: It appears that low values of R lead to better performance, hence we keep $R = 0.1$. This may imply that the reproduction process does create good enough solutions out of the initial population without the need to introduce a lot of diversity across the process;

Hyperparameter	Value	Occurrences
Crossover operator	swap_profiles	0
	swap_orders	0
	swap_weights	0
	mix_criteria_and_weights	18
Mutation operator	random_profile_perturbation	0
	random_weights_perturbation	0
	shrink_profiles	1
	expand_profiles	0
	partially_reverse_order	17
S	100	0
	150	2
	200	1
	250	6
	300	9
B	0.1	1
	0.2	3
	0.3	4
	0.4	6
	0.5	4
R	0.1	10
	0.2	5
	0.3	3
	0.4	0
	0.5	0

Table 1: Number of occurrences of each hyperparameter value among the best configurations in Ω^* . Crossover and mutation operator rows indicate the use of that operator in isolation from the others. The various λ have been grouped, as only one can appear at a time in a configuration. Best choices for each hyperparameter, as well as relevant choices for the operators – according to the observed counts – are highlighted in yellow.

- Crossover and mutation operators: The main conclusion of this process is that some operators clearly outperform the others. Indeed, **mix_criteria_and_weights** and **partially_reverse_order** appear in almost all retained configurations. Consequently, there is no information on the λ parameters.

We remind that all these experiments were performed with pairs of operators taken in isolation from the others. Reported accuracies on Ω^* above show that single-operator settings are sufficient to solve problems already considered complicated in previous literature with high accuracy, in very reasonable time.

5 Numerical analysis

We study in this section the performance of our approach for inferring an SRMP model. We begin by describing the experimental protocol followed by results pertaining to 1. the ability to reproduce comparisons of the ground truth model when P is supposed known; 2. the same ability when no assumption is made on P ; 3. the computation time; 4. the ability to handle noisy data.

Similarly to preliminary experiments in Section 4, we consider here ground truth SRMP models containing $C \in \{7, 11, 15\}$ criteria and $P \in \{1, 2, 3, 4\}$ reference profiles, respecting the constraints introduced in Section 2. We consider a set \mathbf{A} of $A = 500$ random alternatives, and we sample $|\mathcal{D}_{\text{TRAIN}}| \in \{100, 300, 500, 1000, 2000\}$ comparisons. The second set of alternatives \mathbf{A}' consists of $A' = 5000$ new random alternatives, and we build $\mathcal{D}_{\text{TEST}}$ as the set of all comparisons between distinct alternatives in \mathbf{A}' . Finally, since we also want to study the impact on performance of the number of profiles in the inferred solutions compared to the ground truth ones, we denote $\hat{P} \in \{1, 2, 3, 4\}$ the argument given to Algorithm 1, which assumes the real P to be unknown.

5.1 Reproduction of the behavior of a ground truth model

We begin by looking at the ability of our approach to model accurately ground truth SRMP models from a training set. Figure 4 reports the average train and test accuracies when inferring 20 SRMP models with a given number of criteria and profiles.

We observe that in all cases, our approach is able to reach very high values of train accuracy regardless the number of comparisons that were provided. A slight decrease in performance is observed as the amount of input information increases. At the same time, as the number of training comparisons increases, so does the test accuracy, showing a close proximity of the inferred models to the original ones. As the size of the problem (C , number of criteria) increases, the performance slightly degrades. However, the effect of increasing the number of profiles for a given problem size is minimal.

We consider in Figure 5 the average execution time of our algorithm for 20 problem instance sizes and when the SRMP model that needs to be inferred contains between 1 and 4 profiles.

The results show that the execution time increases linearly with the number of provided binary comparisons. A similar but less pronounced trend can also be observed when considering problem instances with more criteria. The impact on the execution time for increasing the number of profiles of the SRMP model is present to an even lesser extent.

5.2 Case when P is not known

We continue by looking at the ability of our approach to construct a model that restores the perspective of the DM when the ground truth number of profiles P is unknown. Figure 6 illustrates the average train accuracy when

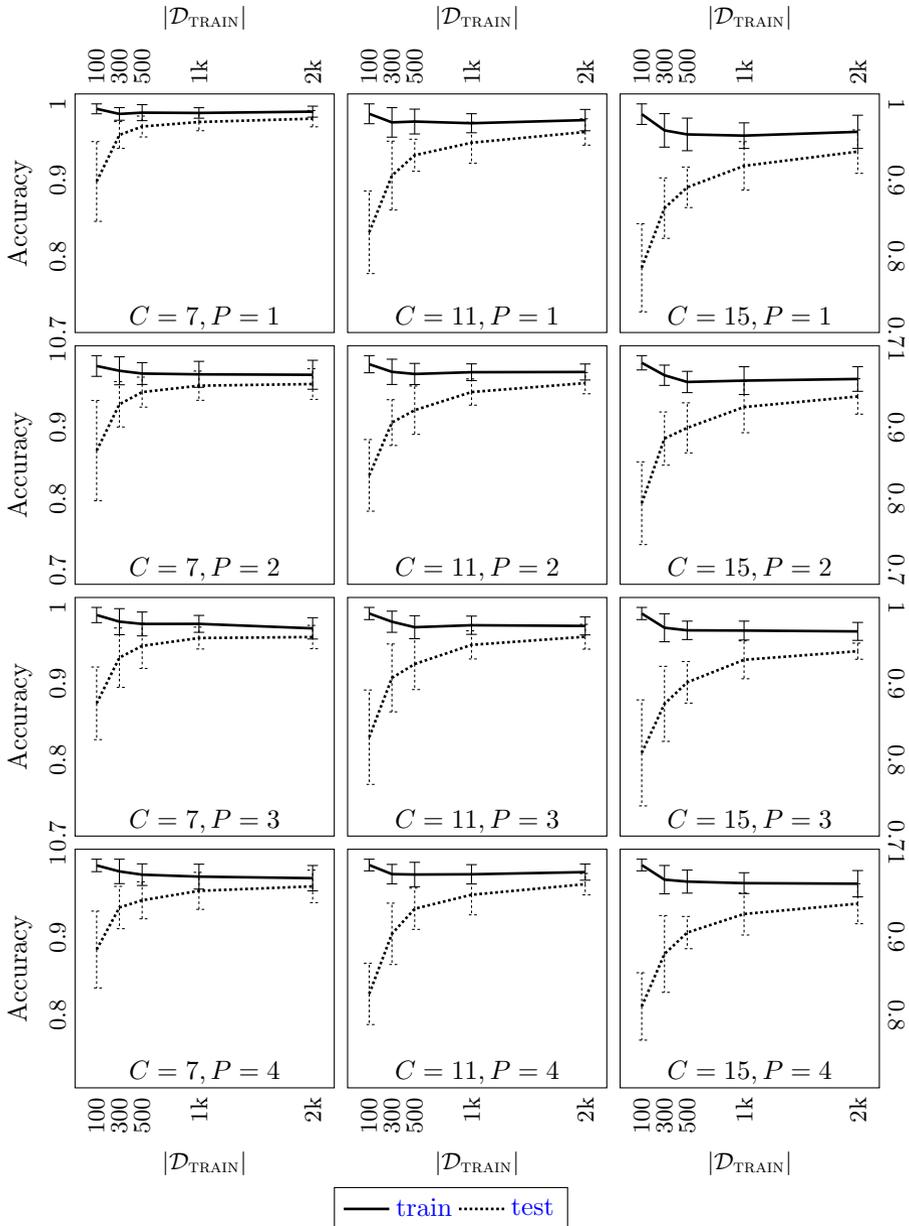


Fig. 4: Average train and test accuracy as a function of the number of training comparisons ($|\mathcal{D}_{\text{TRAIN}}|$) for all combinations of the number of criteria (C) and the number of profiles ($P = \hat{P}$). Bars represent the limits of the 95% confidence intervals.

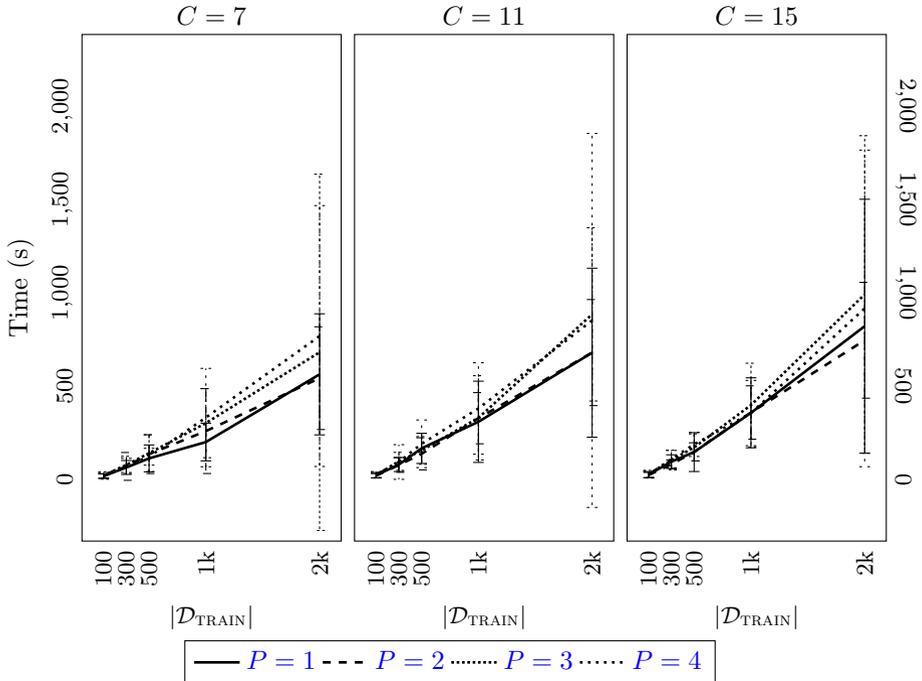


Fig. 5: Average execution time as a function of the number of training comparisons ($|\mathcal{D}_{\text{TRAIN}}|$) for all combinations of the number of criteria (C) and the number of profiles ($P = \hat{P}$). Bars represent the limits of the 95% confidence intervals.

inferring SRMP models with a given number of criteria and a fixed number of profiles. Each point corresponds to the results on 80 problem instances with an equal repartition of the values of P described above.

When $C = 7$ we observe that inferring an SRMP model with a single profile leads to the lowest overall train accuracy as it does not appear to be able to reproduce all of the comparisons produced by underlying models with more profiles. Nevertheless, the train accuracy remains very high, above 0.95. The highest average performance is found when inferring models with two profiles, while, surprisingly, models with more than two profiles perform worse on average despite their potential of modeling more difficult binary comparisons. For more than 7 criteria, the train accuracies remain very high, but the difference between the various models is less apparent.

Figure 7 illustrates the corresponding test accuracy for the previously presented results. For $C = 7$ we observe the same trend where SRMP models with a single profile perform the worst when approximating the perspective of the DM, while models with two profiles perform the best. The difference in performance is somewhat less pronounced for larger numbers of criteria, but reaching

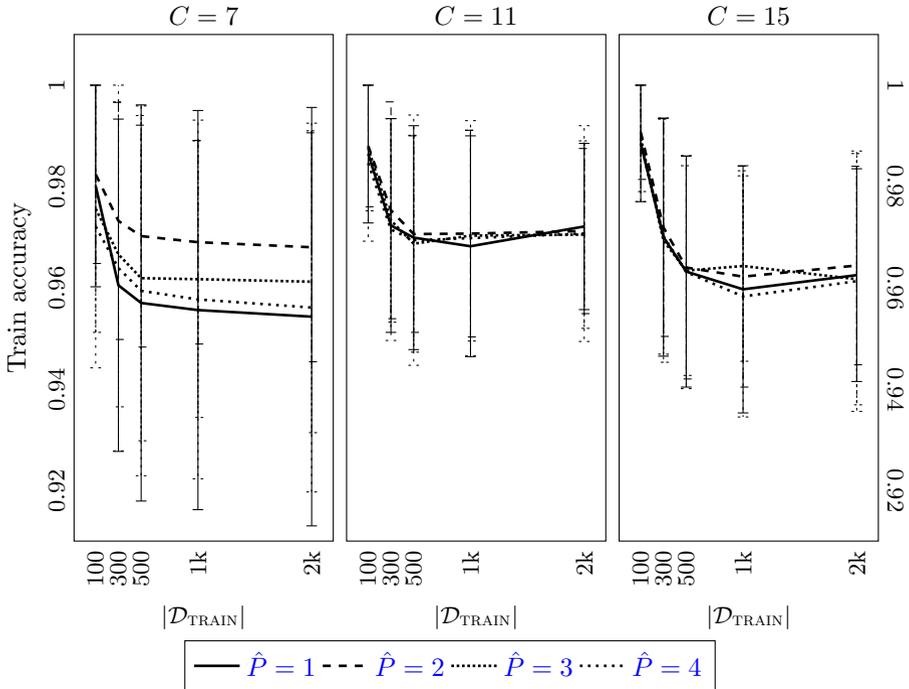


Fig. 6: Average train accuracy as a function of the number of training comparisons ($|\mathcal{D}_{\text{TRAIN}}|$) for any number of criteria (C) when inferring SRMP models with different numbers of profiles (\hat{P}). Bars represent the limits of the 95% confidence intervals.

in all cases very high test accuracy when large sets of binary comparisons are provided.

Looking at the average execution times (Fig 8), we do not observe any significant difference with relation to the results from the previous subsection. The execution times remain linear with relation to the number of provided binary comparisons and increase slightly when more profiles are considered.

5.3 Handling noisy data

Finally, we look at the effect that errors in the learning comparisons may have on the performance of the genetic algorithm as well as the quality of the inferred model. The errors are introduced by randomly selecting a given number of pairwise comparisons in $\mathcal{D}_{\text{TRAIN}}$ and then randomly changing the result of the comparison \odot as follows: if \odot is \succ then \odot becomes either \prec or \sim , if \odot is \prec then \odot becomes either \succ or \sim and if \odot is \sim then \odot becomes either \succ or \prec . In the last case, we use equal probabilities for the two relations, while in the first two cases, the probability of selecting \sim is equal to the proportion of indifference relations in the entire training set. This is motivated by the

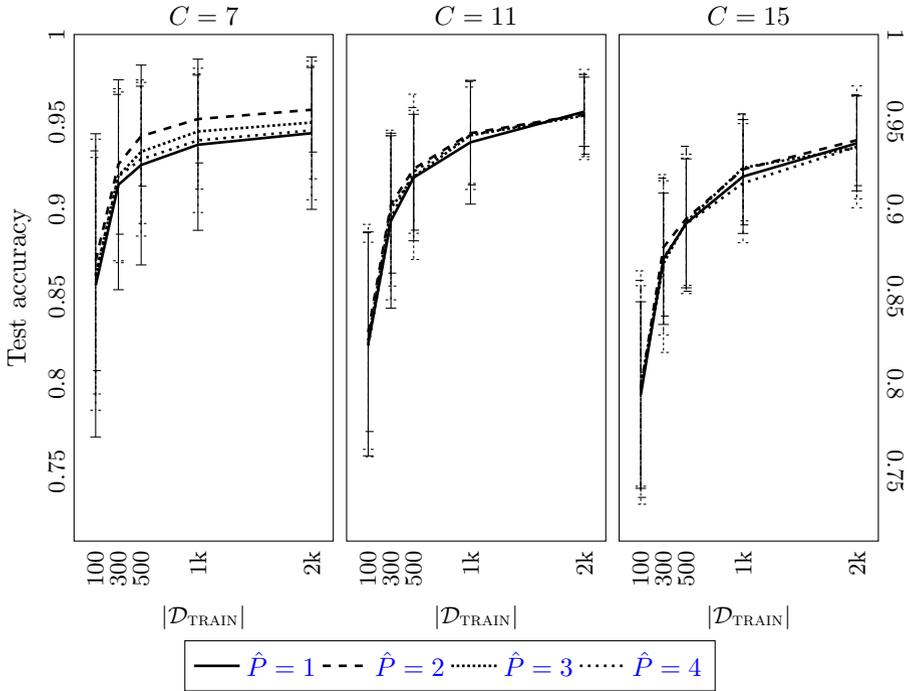


Fig. 7: Average test accuracy as a function of the number of training comparisons ($|\mathcal{D}_{\text{TRAIN}}|$) for any number of criteria (C) when inferring SRMP models with different numbers of profiles (\hat{P}). Bars represent the limits of the 95% confidence intervals.

fact that much fewer pairs of alternatives contain an indifference relation as opposed to a preference relation, and thus we seek to keep this distribution in the training set even when errors are introduced.

We illustrate below only the results when the inferred model matches the number of profiles of the original model, however the same conclusions can be drawn for the other situations.

Figure 9 illustrates the training accuracy over training sets of different sizes, for problem instances containing different numbers of criteria and for SRMP with fewer or more profiles. The lines indicate the training accuracy when different percentages of training comparisons are changed to indicate a different assertion from the one obtained thanks to the original SRMP model. This could either be a preference in the opposite direction or an indifference. Since indifferences between two alternatives have a significantly lower probability of appearing when compared to choices – especially when these alternatives are generated randomly – we make sure to keep the original proportions of indifferences and preferences when inserting errors in the training data.

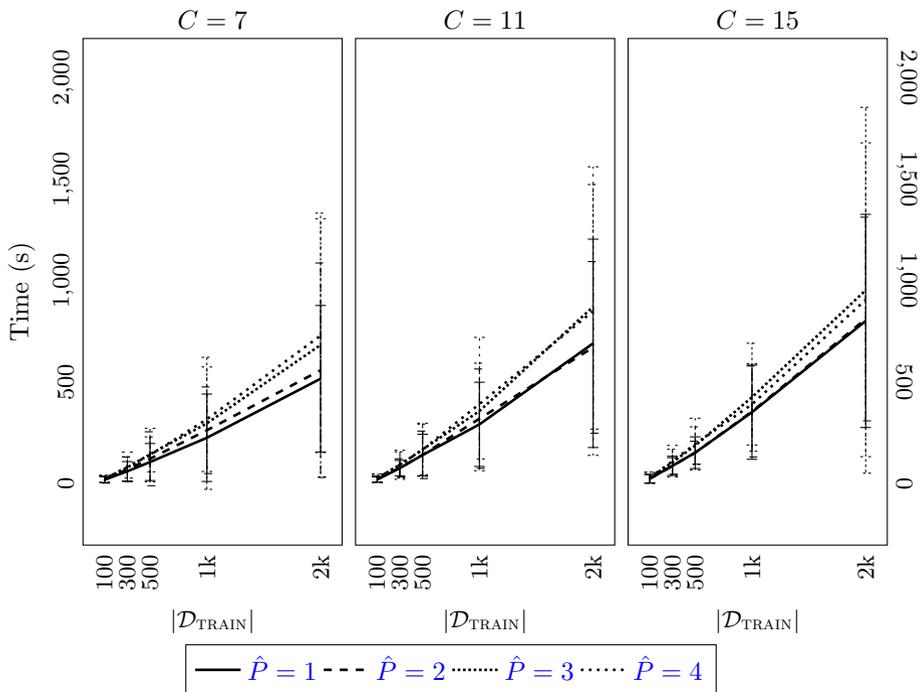


Fig. 8: Average execution time as a function of the number of training comparisons ($|\mathcal{D}_{\text{TRAIN}}|$) for any number of criteria (C) when inferring SRMP models with different numbers of profiles (\hat{P}). Bars represent the limits of the 95% confidence intervals.

We find that, as more and more errors are added to the training data, the training accuracy is significantly impacted. This impact seems to be proportional to the amount of errors, reaching 70% accuracy when up to 30% of the training data is erroneous. For problem instances with more criteria, the effect is less prominent when fewer binary comparisons are considered, however, as the training data increases, so does the penalty on the training accuracy.

Figure 10 illustrates the test accuracy for the same scenarios as above. Despite the rather significant effect on training accuracy, the test accuracy does not seem to suffer the same penalty. This is especially visible for problem instances with fewer criteria. As the number of comparisons in the training data increases, so does the test accuracy. It seems that the genetic algorithm is able to satisfy more of the comparisons that are coherent with the original model and a significant proportion of erroneous comparisons are not fulfilled. Nevertheless, as the number of criteria and the number of reference profiles increase, the penalty incurred by errors in the training data becomes more noticeable.

We do not illustrate the execution time for the previous tests as no difference was observed when introducing errors in the learning data. The algorithm

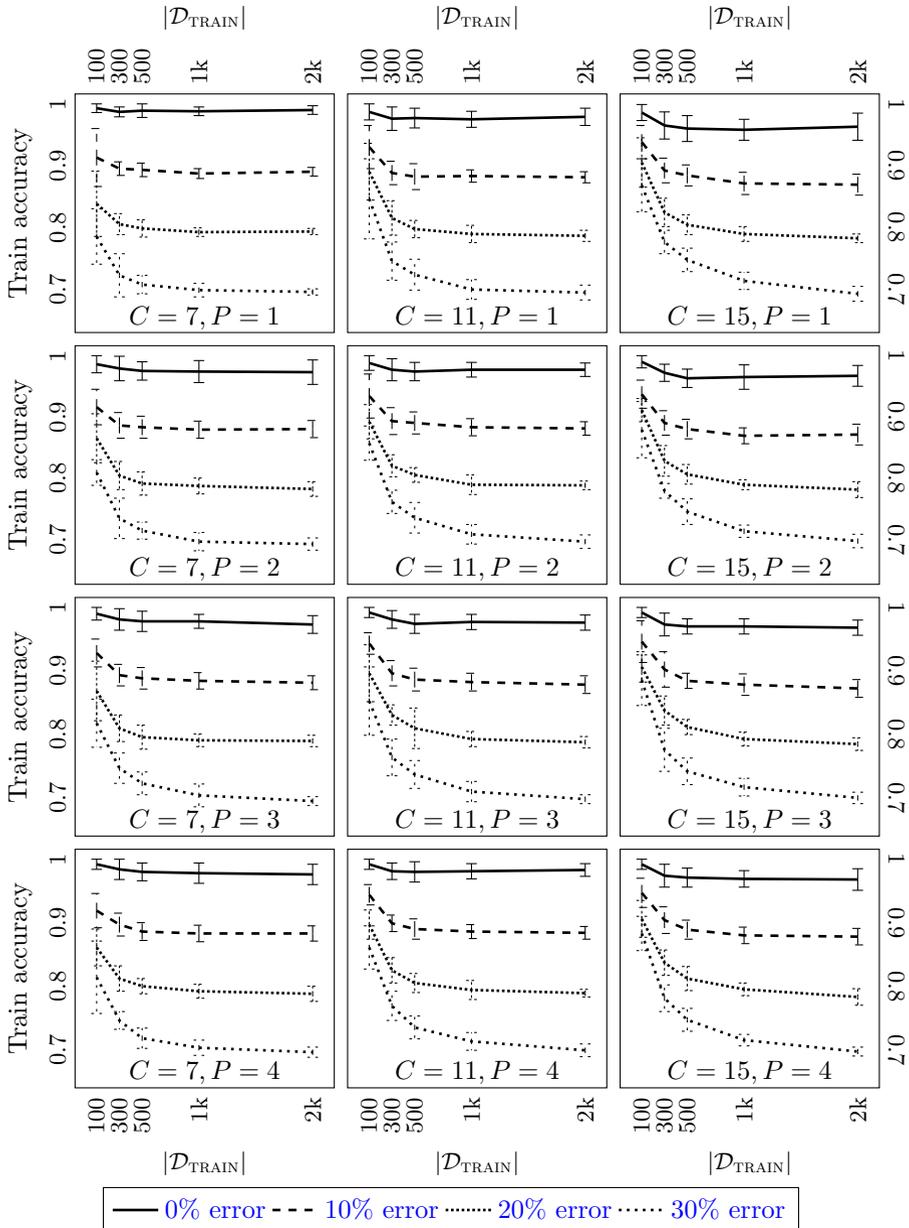


Fig. 9: Average train accuracy as a function of the number of training comparisons ($|\mathcal{D}_{\text{TRAIN}}|$) with relation to the error rate among these comparisons, for all combinations of the number of criteria (C) and the number of profiles in the original and inferred model ($P = \hat{P}$). Bars represent the limits of the 95% confidence intervals.

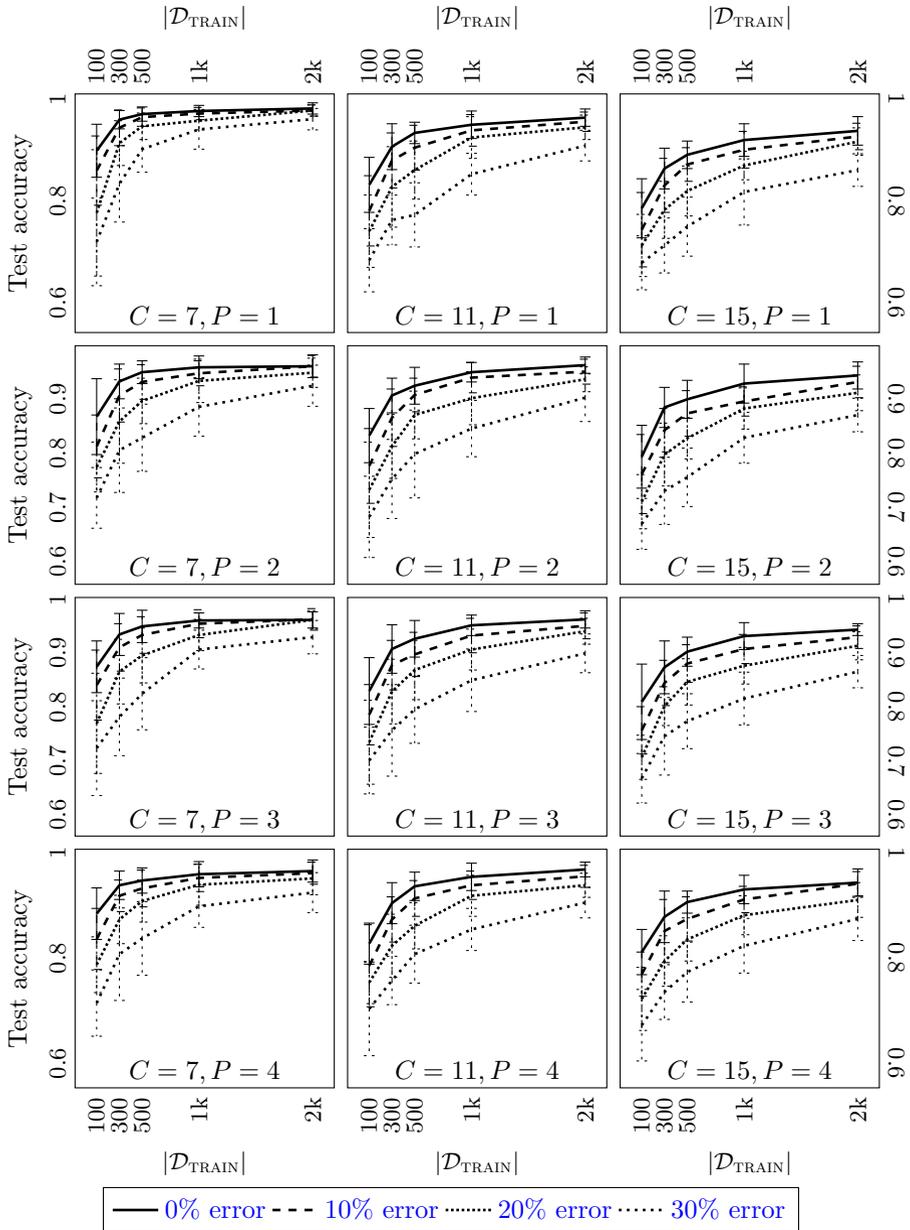


Fig. 10: Average test accuracy as a function of the number of training comparisons ($|\mathcal{D}_{\text{TRAIN}}|$) with relation to the error rate in these comparisons for all combinations of the number of criteria (C) and the number of profiles in the original and inferred model ($P = \hat{P}$). Bars represent the limits of the 95% confidence intervals.

converges in the same amount of time, on average, which is mainly dictated by the number of binary comparisons and the number of criteria.

6 Discussion and conclusion

This article proposes a novel metaheuristic approach for inferring a simple ranking model based on multiple reference profiles. The proposed algorithm is able to handle larger amounts of input information and is more effective in terms of computation times than previous proposals based on mixed-integer linear programming or constraint programming. The extensive experiments conducted in this study demonstrate a generalization power of the learned models similar to the ones reported in previous works, but in significantly lower computation times, even in presence of noisy training comparisons. This allows the proposed approach to be used in real-life problems involving many criteria.

Table 2 illustrates the largest problem instances handled by previous works, in terms of number of criteria (C), number of reference profiles (P), number of input binary comparisons ($|\mathcal{D}_{\text{TRAIN}}|$). In some cases, two problem configurations are associated to a single reference. The results reported by these works, in terms of training and testing accuracy as well as computation time, are given in middle three lines of the table. We have also tested the approach we propose on all problem configurations, averaging the results over 10 random DMs and constructing the training set using a set of 50 randomly generated alternatives. These results are given in the final three lines of the table.

State of the art reference		[17]	[17]	[3]	[3]	[16]
$C =$		5	7	5	6	10
$P =$		2	2	3	2	3
$ \mathcal{D}_{\text{TRAIN}} =$		100	100	1000	600	500
State of the art results	Training acc.	1	1	1	1	0.97
	Testing acc.	0.89	0.86	0.98	0.97	n/a
	Time	6m16s	33m50s	8m45s	18m20s	3m2s
Proposed approach results	Training acc.	0.988	0.987	0.982	0.977	0.975
	Testing acc.	0.865	0.832	0.937	0.913	0.875
	Time	16.852s	17.86s	4m39s	2m37	2m27s

Table 2: Comparison of the largest problem instances handled by previously published works with ours, averaged for 10 random DMs and choices of A of two different orders, using the settings described in Section 4.

When the time budget is not a limitation and problems are small, approaches based on exact methods obviously provide better training accuracy than our metaheuristic approach. However, training accuracy of our algorithm remains very high in all considered settings, and the required time is generally a lot smaller. This allows us to tackle larger problems than previous works.

This is especially true given that our implementation is done in Python without any optimization, while [17] and [16] make use of the Cplex solver. [3] makes use of solvers CryptoMiniSAT and Maxino. A C/C++ implementation of our approach would therefore lead to higher gains for those interested in this aspect.

Complementing Table 2, we have observed that our approach can solve problems with $C = 15$, $P = 4$, $A = 50$ and $|\mathcal{D}_{\text{TRAIN}}| = 2000$ with high accuracy in about 19m. Therefore, for more or less the same executing time budget, our approach is able to handle problem instances three times larger in terms of the number of criteria and 20 times larger in terms of the number of input binary comparisons as opposed to a mixed-integer linear programming approach.

Still, a certain amount of questions and perspectives arise from our work. The performance of the proposed approach could probably be further improved by incorporating various problem-specific features, such as heuristics, local search, or hybridization with other optimization techniques. An adaptation for inferring the number of reference profiles without requiring this information as an input parameter could also be done. Last but not least, a multi-operator setting with multiple crossover and mutation operators could also be explored to improve the performance of the algorithm. They could either be called randomly, or depending on some past performances through a reinforcement approach.

Also, the proposed genetic algorithm handles a population of solutions, in which currently only one solution (the best in terms of training accuracy) is output. It could be worth looking in further details at the other “good” solutions of the population, to determine if some of their characteristics could be used to improve the best element of the population.

Finally, a limitation of all approaches in the literature – including ours – is the absence of focus on generalization. Indeed, one may be interested not particularly in the ability of the inferred model to reproduce existing decisions, but more to mimic the behavior a DM would have in new unseen situations. Drawing inspiration from the machine learning field, we believe that introducing a validation set, extracted from the training comparisons, to serve as a proxy for future testing data, would be worth exploring to maximize the generalization ability of inferred SRMP models. Experiments in this work still hint that maximizing training accuracy implies an increase in testing accuracy, and it seems that there is no overfitting problem here (as we are learning a model and not a classification/regression function). Still, taking generalization ability into account during model inference may help convergence. This is in our opinion a great direction for future work.

Conflict of Interest: The authors declare that they have no conflict of interest.

Data availability: The datasets generated during and/or analyzed during the current study are available in the GitHub repository, https://github.com/BastienPadeloup/learn_srmp.

References

- [1] J. Almeida Dias, J. Figueira, and B. Roy. ELECTRE TRI-C: A multiple criteria sorting method based on characteristic reference actions. *European Journal of Operational Research*, 204(3):565–580, 2010.
- [2] C.A. Bana e Costa and J-C. Vansnick. MACBETH—An interactive path towards the construction of cardinal value functions. *International transactions in operational Research*, 1(4):489–500, 1994.
- [3] K. Belahcene, V. Mousseau, W. Ouerdane, M. Pirlot, and O. Sobrie. Ranking with multiple reference points: Efficient SAT-based learning procedures. *Computers & Operations Research*, 150:106054, 11 2022.
- [4] K. Belahcene, V. Mousseau, W. Ouerdane, M. Pirlot, and O. Sobrie. Multiple criteria sorting models and methods-Part I: survey of the literature. *4OR: A Quarterly Journal of Operations Research*, January 2023.
- [5] K. Belahcene, V. Mousseau, W. Ouerdane, M. Pirlot, and O. Sobrie. Multiple criteria sorting models and methods. Part II: theoretical results and general issues. *4OR: A Quarterly Journal of Operations Research*, January 2023.
- [6] D. Bouyssou and T. Marchant. Multiattribute preference models with reference points. *European Journal of Operational Research*, 229(2):470–481, 2013.
- [7] D. Bouyssou, T. Marchant, M. Pirlot, A. Tsoukiàs, and P. Vincke. *Evaluation and decision models with multiple criteria: Stepping stones for the analyst*. International Series in Operations Research and Management Science, Volume 86. Boston, 1st edition, 2006.
- [8] J. Butler, J. Jia, and J. Dyer. Simulation techniques for the sensitivity analysis of multi-criteria decision models. *European Journal of Operational Research*, 103(3):531–546, 1997.
- [9] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI, 1975. second edition, 1992.
- [10] C-L. Hwang, Y-J. Lai, and T-Y. Liu. A new approach for multiple objective decision making. *Computers & Operations Research*, 20(8):889–899, 1993.
- [11] E. Jacquet-Lagrèze and Y. Siskos. Preference disaggregation: 20 years of MCDA experience. *European Journal of Operational Research*, 130(2):233–245, 2001.
- [12] A. Khannoussi, N. Benabbou, A.-L. Olteanu, and P. Meyer. Incremental elicitation of the criteria weights of srmp using a regret-based query selection strategy. In *DA2PL 2020: From Multiple Criteria Decision Aid to Preference Learning*, 2020.
- [13] A. Khannoussi, A.-L. Olteanu, C. Dezan, J.-P. Diguët, C. Labreuche, J. Petit-Frère, and P. Meyer. Incremental learning of simple ranking method using reference profiles models. In *DA2PL 2018: From Multiple Criteria Decision Aid to Preference Learning*, 2018.
- [14] A. Khannoussi, A.-L. Olteanu, C. Labreuche, and P. Meyer. Simple ranking method using reference profiles: incremental elicitation of the

- preference parameters. *JOR: A Quarterly Journal of Operations Research*, 20(3):499–530, 2022.
- [15] A. Leroy, V. Mousseau, and M. Pirlot. Learning the parameters of a multiple criteria sorting method. In R. Brafman, F. Roberts, and A. Tsoukiàs, editors, *Algorithmic Decision Theory*, volume 6992, pages 219–233. Springer, 2011.
- [16] J. Liu, W. Ouerdane, and V. Mousseau. A metaheuristic approach for preference learning in multicriteria ranking based on reference points. in Proceeding of the 2nd workshop “From Multiple Criteria Decision Aid to Preference Learning (DA2PL)”, 2014.
- [17] A-L. Olteanu, K. Belahcène, V. Mousseau, W. Ouerdane, A. Rolland, and J. Zheng. Preference elicitation for a ranking method based on multiple reference profiles. *JOR: A Quarterly Journal of Operations Research*, 20(1):63–84, 2022.
- [18] A. Rolland. Reference-based preferences aggregation procedures in multicriteria decision making. *European Journal of Operational Research*, 225(3):479–486, 2013.
- [19] B. Roy. The outranking approach and the foundations of ELECTRE methods. *Theory and Decision*, 31:49–73, 1991.
- [20] B. Roy. *Multicriteria Methodology for Decision Aiding*. Kluwer Academic, Dordrecht, 1996.
- [21] A. Tversky and D. Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, 106(4):1039–1061, 11 1991.
- [22] D. von Winterfeldt and W. Edwards. *Decision Analysis and Behavioral Research*. Cambridge University Press, 1986.