Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Combining Preference Elicitation and Search in Multiobjective State-Space Graphs

Nawal Benabbou and **Patrice Perny**
Sorbonne Universites, UPMC Univ Paris 06, UMR 7606, LIP6
CNRS, UMR 7606, LIP6, F-75005, Paris, France
4 Place Jussieu, 75005 Paris, France
nawal.benabbou@lip6.fr, patrice.perny@lip6.fr

Abstract

The aim of this paper is to propose a new approach interweaving preference elicitation and search to solve multiobjective optimization problems. We present an interactive search procedure directed by an aggregation function, possibly non-linear (e.g. an additive disutility function, a Choquet integral), defining the overall cost of solutions. This function is parameterized by weights that are initially unknown. Hence, we insert comparison queries in the search process to obtain useful preference information that will progressively reduce the uncertainty attached to weights. The process terminates by recommending a near-optimal solution ensuring that the gap to optimality is below the desired threshold. Our approach is tested on multiobjective state space search problems and appears to be quite efficient both in terms of number of queries and solution times.

1 Introduction

In many practical search problems considered in Artificial Intelligence (e.g path planning, game search, preference-based configuration), the evaluation and comparison of solutions involve several aspects or points of view (e.g. in path planning, time, distance, energy consumption, risk). For this reason, standard search algorithms are worth generalizing to be implementable in the context of multiobjective optimization. This statement has motivated various contributions stemming from the initial A^* search algorithm [Hart *et al.*, 1968] and aiming at proposing extensions to cope with multiple conflicting criteria. Let us mention, among others, MOA* the multiobjective extensions of A* finding all Pareto optimal cost vectors [Stewart and White III, 1991; Mandow and de la Cruz, 2005] in a vector-valued state space graph, U* a variation of MOA* used to find a path maximizing a multiattribute utility function [Dasgupta *et al.*, 1995], and a preference-based specialization of MOA* [Perny and Spanjaard, 2003]. The same trend can be observed for AND/OR search [Dasgupta *et al.*, 1996a; Marinescu, 2010], game search [Dasgupta *et al.*, 1996b] and constraint optimization [Marinescu *et al.*, 2013].

In preference-based search, the preference model is often assumed to be known and the effort is put on algorithmic issues. Thus, the elicitation problem must be solved in a prior stage. The standard elicitation procedures proposed in multiattribute utility theory aim at providing a complete elicitation [Fishburn, 1967; Krantz *et al.*, 1971; Keeney and Raiffa, 1976]; the preference model is precisely constructed on the entire multiattribute space. This approach is however difficult to implement on combinatorial domains, except perhaps for very simple and decomposable utility models. The development of recommender systems and the need of fast and efficient preference elicitation procedures for large databases have led researchers to propose less ambitious elicitation procedures: one seeks to obtain only a part of the preference model, sufficient to make a decision on the given instance. This suggests resorting to more incremental elicitation processes. Preference queries are selected one at a time, to be as informative as possible, so as to progressively reduce the set of admissible utility functions until a robust decision can be made. In this line, let us mention the ISMAUT method [White III *et al.*, 1984] for the elicitation of multiattribute utility functions, and more recently, strategies developed within the artificial intelligence community for preference query selection using the minimax-regret criterion, see e.g., [Boutilier, 2002; Wang and Boutilier, 2003; Boutilier *et al.*, 2006; Lu and Boutilier, 2011; 2013].

Incremental elicitation procedures are also involved in the context of voting with partial preference profiles. When individual preferences are incomplete, one can study possible and necessary winners (e.g., [Konczak and Lang, 2005; Xia and Conitzer, 2011; Lang *et al.*, 2012; Ding and Lin, 2013]). In this setting, incremental elicitation methods are used to progressively reduce the set of possible winners until a winner can be determined with some guarantee [Kalech *et al.*, 2010; Lu and Boutilier, 2011; Dery *et al.*, 2014]). The elicitation task consists in obtaining new individual preference judgements over candidates given explicitly. In this paper, we consider a slightly different elicitation context. Our aim is to resort to incremental preference elicitation to refine a multiobjective state-space search procedure. One-dimensional preferences are assumed to be known and represented by criterion functions. The elicitation burden is focused on the determination of weights used in the aggregation phase to define overall preferences over a combinatorial set of alternatives (implicitly defined as the solution paths of a state space graph).

Incremental elicitation strategies based on minimax-regret minimization have proven quite effective but require to minimize regrets for every pair of feasible solutions. This fits well to decision problems on explicit sets of alternatives. This can be used for multiobjective combinatorial optimization problems as well when the set of Pareto-optimal alternatives is not too large and can be computed in a first stage prior to preference elicitation. Our aim here is to propose a more direct approach that consists in interweaving preference elicitation and search. Starting with an initial set of possible utility functions characterized by some weighting parameters, we propose to generate preference queries during the search so as to progressively reduce the set of possible weights until an optimal solution can be determined or approximated with some guarantees. In this process, the decision model is progressively revealed and constructed during the search. However, in general, a robust solution can be found without completely specifying the model. We want to apply and test this approach on two classes of utility models. We consider first *additive utility functions* [Fishburn, 1968] parameterized by weights representing the importance of attributes. Then we will consider a more general model, namely the *Choquet Expected Utility* [Schmeidler, 1986] parameterized by a set function defining the importance of all coalitions of attributes.

A first attempt in this direction has been recently proposed for linear weighted aggregators (which are a special case of additive utilities) [Benabbou and Perny, 2015]. However, it does not extend to non-linear multiattribute utility functions because the proposed algorithm relies on pruning rules based on the Bellman principle. Unfortunately this principle does not hold anymore when multiobjective costs of paths are aggregated with a non-linear function. Another recent study concerns the case of incremental elicitation of capacity weights in Choquet integrals (see [Benabbou *et al.*, 2014]) but assumes that the set of alternatives is given explicitly. Here we propose an approach to overcome both difficulties simultaneously: the non-linearity of the multiattribute utility function and the combinatorial nature of the set of alternatives.

The paper is organized as follows: in Section 2, we introduce the formal framework and recall some background on decision models and preference-based search. Then, Section 3 and 4 are devoted to the introduction of our procedure combining elicitation and search. The efficiency of this approach will be discussed in Section 5 where numerical experiments are reported to assess the performance of the search procedure both in terms of number of queries and computation times.

2 Preference-based Search in MO Graphs

We consider $G=(N,A)$ a state space graph where N denotes the finite set of nodes representing all states and A is the set of arcs representing the admissible transitions. Formally, $A=\{(n,n'): n\in N, n'\in S(n)\}$ where $S(n)\subseteq N$ is the set of all nodes that can be reached from node n by a single transition. The set of all paths between node n and node n' is denoted $P(n,n')$, and each of them is characterized by a list of nodes of type $\langle n,\dots,n'\rangle$. In particular, the set of solution paths, starting at source node $s\in N$ and reaching any goal node $\gamma\in\Gamma$, is denoted $P(s,\Gamma)$. Besides, we consider $Q=\{1,\dots,q\}$ a finite set of criteria (e.g. time, distance) represented by q cost functions $g_i: A\rightarrow\mathbb{R}_+$, $i\in Q$. Hence, each path p in graph G is associated with a cost vector denoted $g(p)=(g_1(p),\dots,g_q(p))$ and $g_i(p)=\sum_{(n,n')\in p} g_i((n,n'))$ for all $i\in Q$. Finally, the image of all solution paths in the space of criteria is $\mathcal{X}=\{g(p): p\in P(s,\Gamma)\}$ and $\mathrm{ND}(\mathcal{X})$ is set of Pareto-optimal vectors in $\mathcal{X}$.

Since we are in a context of cost minimization, we will use disutility functions to be minimized rather than utility functions to be maximized. Hence, we consider a multiattribute disutility function $\psi_\omega^v:\mathbb{R}_+^q\rightarrow[0,1]$ which associates the disutility: $\psi_\omega^v(x)=\psi_\omega(v_1(x_1),\dots,v_q(x_q))$ to any cost vector $x=(x_1,\dots,x_q)$, where $v_i:\mathbb{R}_+\rightarrow[0,1]$ is a disutility function measuring the subjective cost of consequence x_i for the Decision Maker (DM) , and $\psi_\omega:[0,1]^q\rightarrow[0,1]$ is a scalarizing function with a parameter denoted ω. A solution x is preferred to another solution y when $\psi_\omega^v(x)\le\psi_\omega^v(y)$. We consider here that one-dimensional disutility functions v_i have been already elicited in a prior step, using standard techniques (see e.g [Keeney and Raiffa, 1976; Bana e Costa and Vansnick, 2000]) and we focus on the elicitation of paremeters $\omega=(\omega_1,\dots,\omega_m)$ so as to approximate DM's preferences with a proper scalarizing function. Throughout the paper, we will consider two main families of utility functions:

- *Additive utilities:* $U_\omega^v(x)=\sum_{i=1}^q\omega_i v_i(x_i)$ where $\omega=(\omega_1,\dots,\omega_q)$ is a vector of positive weights adding up to 1. Note that this family virtually includes also quasi-arithmetic means of the form $M_\omega^\phi(x)=\phi^{-1}(\sum_{i=1}^q\omega_i\phi(x_i))$ for strictly increasing ϕ. Minimizing $M_\omega^\phi(x)$ is indeed equivalent to minimizing $\phi(M_\omega^\phi(x))=U_\omega^v(x)$ for $v_i(z)=\phi(z)$. This includes as special case the weighted L_p norm obtained for $\phi(z)=z^p$, the geometric mean for $\phi(z)=\ln(z)$ and the standard weighted sum for $\phi(z)=z$. More generally additive utilities can use distinct one-dimensional disutility functions v_i to encode preferences on the different criteria.

- *Choquet expected utilities:* generalize additive utilities by $C_\omega^v(x)=\sum_{i=1}^q\left[v_{(i)}(x_{(i)})-v_{(i-1)}(x_{(i-1)})\right]\omega(X_{(i)})$ where $(.)$ denotes a permutation of $(1,\dots,q)$ such that $v_{(i)}(x_{(i)})\le v_{(i+1)}(x_{(i+1)})$ for all $i=1,\dots,q-1$, $x_{(0)}$ is a ficticious value such that $v_i(x_{(0)})=0$, and $X_{(i)}=\{(i),\dots,(q)\}$ is the subset of indices $j\in Q$ corresponding to the $n+1-i$ largest disutility values $v_j(x_j)$. In this case ω is a *capacity*, i.e. a set-function $\omega: 2^Q\rightarrow[0,1]$ such that $\omega(\emptyset)=0$, $\omega(Q)=1$, and $\omega(A)\le\omega(B)$ whenever $A\subseteq B$ (monotonicity). The monotonicity condition ensures that $C_\omega^v(x)\le C_\omega^v(y)$ whenever x Pareto-dominates y (denoted $x\prec_P y$ hereafter), i.e. $x_i\le y_i$ for all $i\in Q$ and $x_j<y_j$ for some $j\in Q$. The value $\omega(X)$ represents the importance attached to a coalition of criteria $X\subseteq Q$. When $\omega(X)=\sum_{i\in X}\omega(\{i\})$, ω is said to be additive and C_ω^v boils down to U_ω^v. For more details see [Grabisch *et al.*, 2009].

Example 1. *Assume we have 3 criteria. Let ω be a capacity defined on 2^Q, for $Q=\{1,2,3\}$, as follows:*

	$\{1\}$	$\{2\}$	$\{3\}$	$\{1,2\}$	$\{1,3\}$	$\{2,3\}$
ω	0.4	0.6	0.5	0.8	0.7	0.9

If $x = (14, 12, 10)$, $y = (8, 16, 12)$ and $u_i(z) = (z/20)^2$ for all $i \in Q$, we obtain that x is preferred to y because:

$$\psi^v_\omega(x) = .5^2 + (.6^2 - .5^2)\omega(\{1,2\}) + (.7^2 - .6^2)\omega(\{1\}) = .4$$
$$\psi^v_\omega(y) = .4^2 + (.6^2 - .4^2)\omega(\{2,3\}) + (.8^2 - .6^2)\omega(\{2\}) = .5$$

We wish to emphasize here the interest, from a descriptive and prescriptive viewpoint, of resorting to non-linear multiattribute utility functions. This not only provides a more general and flexible class of decision models that can be tuned to the observed preferences, but it also enables to enhance the possibility of finding good compromise solutions within the Pareto set. Let us consider indeed the example given in Figure 1 based on a biobjective shortest path problem. In the figure, every point represents a feasible cost vector and red points represent the Pareto set. As can be seen from the convex hull of these points, only three points in the Pareto set can be obtained by minimizing a weighted sum of the costs. When minimizing a convex utility U^v_ω (e.g. $v_i(x) = x^2$, Fig.1 left) or a convex Choquet integral C^v_ω (e.g. $v_i(x) = x$, Fig.1 right), we can see from the isopreference curves plotted in Figure 1 that more interesting compromise solutions can be found, even when they do not belong to the boundary of the convex hull of the feasible points.

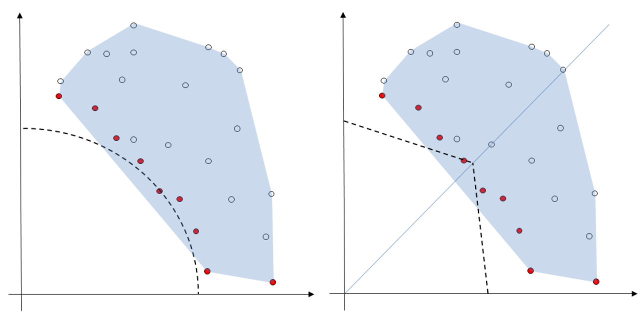

Figure 1: Optimum with U^v_ω and C^v_ω

Note that, whenever $\psi^v_\omega = U^v_\omega$ or $\psi^v_\omega = C^v_\omega$, the inequality $\psi^v_\omega(x) \leq \psi^v_\omega(y)$ is linear in ω for any fixed cost vectors $x, y \in \mathbb{R}^q_+$. Hence, any preference judgement of type "x is preferred to y" will be translated as a linear constraint bounding the set of admissible weighting vectors Ω. Therefore, when preference judgements are obtained from the DM, the set of admissible weights Ω is restricted by linear constraints and thus is a convex polyhedron. This will be useful for performing optimization with imprecise parameters ω.

Robust recommendations with minimax regrets

As ω is imprecisely known, a solution which remains ψ^v_ω-optimal for all $\omega \in \Omega$ may not exist. We face a decision problem under uncertainty where Ω is the set of states of nature and any cost vector x is associated with an act (according to the definition of Savage [Savage, 1954]) characterized by the set of consequences $\{\psi^v_\omega(x), \omega \in \Omega\}$. In this context, we are concerned with the determination of a robust solution, i.e. a flexible solution preserving nice perspectives with respect to the possible future evolutions of the uncertainty set Ω. More precisely, the robust solutions can be defined as those minimizing the max-regret criterion [Wang and Boutilier, 2003]. They are characterized by the following definitions, for all $x, y \in \mathcal{X}$:

Pairwise Max Regret: $\mathrm{PMR}(x, y, \Omega) = \max_{\omega \in \Omega}\{\psi^v_\omega(x) - \psi^v_\omega(y)\}$

Max Regret: $\mathrm{MR}(x, \mathcal{X}, \Omega) = \max_{y \in \mathcal{X}} \mathrm{PMR}(x, y, \Omega)$

Minimax Regret: $\mathrm{MMR}(\mathcal{X}, \Omega) = \min_{x \in \mathcal{X}} \mathrm{MR}(x, \mathcal{X}, \Omega)$

$\mathrm{MR}(x, \mathcal{X}, \Omega)$ is the worst-case regret of choosing x instead of any $y \in \mathcal{X}$. Robust solutions are those minimizing MR values over $\mathcal{X}$. However, given the set Ω, the worst-case loss measured by MMR might be too large for certifying the quality of the solution. In this case we are going to collect new preference information so as to reduce the uncertainty set Ω and therefore the MMR. Note that $x \prec_P y$ implies that $\mathrm{PMR}(z, x, \Omega) \geq \mathrm{PMR}(z, y, \Omega)$ and $\mathrm{PMR}(x, z, \Omega) \leq \mathrm{PMR}(y, z, \Omega)$ for any solution $z \in \mathcal{X}$. Hence, Pareto-dominated solutions can be omitted during the search since $\mathrm{MMR}(\mathcal{X}, \Omega) = \mathrm{MMR}(\mathrm{ND}(\mathcal{X}), \Omega)$.

3 Search with Imprecise Parameters

We introduce now a general interactive elicitation procedure alternating preference elicitation steps and search steps. The search steps of the procedure are based on recent variants of MOA* [Mandow and de la Cruz, 2005] and U^* [Dasgupta *et al.*, 1995; Perny and Spanjaard, 2003], adapted to minimize regrets under utility uncertainty. Let us recall now the standard concepts and formalism for multiobjective state-space search. In vector-valued graphs, there possibly exists several optimal paths with different cost vectors to reach a given node. Therefore, the basic graph exploration procedure consists in iteratively expanding labels attached to subpaths rather than nodes. Labels are of the form $\ell = [n_\ell, p_\ell, g_\ell]$ where p_ℓ denotes a path from s to n_ℓ and $g_\ell = g(p_\ell)$ denotes its cost. At any iteration of the algorithm, a label is selected for expansion. The expansion of a label ℓ^* generates the set of its successors $\{[n, p_{\ell^*} \circ n, g(p_{\ell^*} \circ n)] : n \in S(n_{\ell^*})\}$. The set of generated labels is divided into two disjoint sets: a set $\mathcal{C}$ of *closed* labels (yet expanded) and a set $\mathcal{O}$ of *open* labels (candidate to expansion). The set $\mathcal{C}$ (resp. $\mathcal{O}$) restricted to labels ℓ such that $p_\ell \in P(s, n)$ is denoted $\mathcal{C}(n)$ (resp. $\mathcal{O}(n)$). Moreover, the expanded labels corresponding to the current possibly optimal solution paths are stored in a set denoted $\mathcal{S}$ and the corresponding set of cost vectors is denoted $g_{\mathcal{S}}$. Another feature imported from MOA* is that, for each generated label ℓ, a set $F(\ell) = \{g_\ell + h : h \in H(n_\ell)\}$ of cost vectors is computed to estimate the cost vectors of the solution paths extending p_ℓ, where $H(n_\ell)$ is a set of heuristic costs estimating the set $\{g(p) : p \in P(n_\ell, \Gamma)\}$.

We consider now the problem of finding an optimal solution path for the minimax regret decision criterion such that the gap to optimality, quantified by the minimax regret MMR, is bounded above by threshold δ. First, we propose a pruning rule that enables, given a set of feasible weights Ω, to detect subpaths that necessarily lead to solutions with a max regret MR strictly greater than δ. This rule is based on the following dominance relation.

Definition 1 ($\lhd^\delta_\Omega$-dominance). *For all $X, Y \subset \mathbb{R}^q_+$:*

$$X \lhd^\delta_\Omega Y \;\Leftrightarrow\; \forall y \in Y,\, \forall \omega \in \Omega,\, \exists x \in X,\; \psi^v_\omega(y) - \psi^v_\omega(x) > \delta$$

Then, the following property holds.

Proposition 1. *For all $X, Y \subset \mathbb{R}^q_+$:*

$$X \lhd^\delta_\Omega Y \;\Leftrightarrow\; \forall y \in Y, \min_{\omega\in\Omega}, \max_{x\in X} \left[\psi^v_\omega(y) - \psi^v_\omega(x)\right] > \delta$$

Proof. Consider $X, Y \subset \mathbb{R}^q_+$ such that $X \lhd^\delta_\Omega Y$ and let $y \in Y$. Then for all $\omega \in \Omega$, there exists $x \in X$ such that $\psi^v_\omega(y) - \psi^v_\omega(x) > \delta$. Therefore, for all $\omega \in \Omega$, we have $\max_{x\in X}[\psi^v_\omega(y) - \psi^v_\omega(x)] > \delta$, and in particular we have $\min_{\omega\in\Omega}\max_{x\in X}[\psi^v_\omega(y) - \psi^v_\omega(x)] > \delta$. Consider now $X, Y \subset \mathbb{R}^q_+$ such that $\min_{\omega\in\Omega}, \max_{x\in X}[\psi^v_\omega(y) - \psi^v_\omega(x)] > \delta$ for all $y \in Y$. In that case, for all $y \in Y$ and all $\omega \in \Omega$, we have $\max_{x\in X}[\psi^v_\omega(y) - \psi^v_\omega(x)] > \delta$ and so there exists $x \in X$ such that $\psi^v_\omega(y) - \psi^v_\omega(x) > \delta$. Hence we have $X \lhd^\delta_\Omega Y$. □

Thus, since Ω is a convex polyhedron and $\psi^v_\omega(x)$ is linear in ω for any fixed $x \in \mathbb{R}^q_+$, $\lhd^\delta_\Omega$-dominance tests can efficiently be performed using linear programming. Hence, we propose a pruning rule based on the following proposition:

Proposition 2. *For any $\ell' \in \mathcal{O}$, if $g_\mathcal{S} \lhd^\delta_\Omega F(\ell')$, then path $p_{\ell'}$ cannot be part of a solution path with a* MR *below δ.*

Proof. Let $\ell' \in \mathcal{O}$ be such that $g_\mathcal{S} \lhd^\delta_\Omega F(\ell')$. For any path $p' \in P(n_{\ell'}, \Gamma)$ and any $\Omega' \subseteq \Omega$, we want to prove that $\mathrm{MR}(g(p_{\ell'} \circ p'), \mathcal{X}, \Omega') > \delta$. Since H is admissible, there exists $h' \in H(n_{\ell'})$ such that h' Pareto-dominates $g(p')$, and so $g_{\ell'} + h'$ Pareto-dominates $g_{\ell'} + g(p') = g(p_{\ell'} \circ p')$. Since ψ^v_ω is increasing with Pareto-dominance, then we have $\psi^v_\omega(g_{\ell'} + h') \leq \psi^v_\omega(g(p_{\ell'} \circ p'))$ for all $\omega \in \Omega'$. Moreover, since we have $g_\mathcal{S} \lhd^\delta_\Omega F(\ell') = \{g_{\ell'} + h : h \in H(n_{\ell'})\}$, then for all $\omega \in \Omega'$, there exists $\ell \in \mathcal{S}$ such that $\psi^v_\omega(g_{\ell'} + h') - \psi^v_\omega(g_\ell) > \delta$, and so we have $\psi^v_\omega(g(p_{\ell'} \circ p')) - \psi^v_\omega(g_\ell) > \delta$. Hence, for all $\omega \in \Omega'$, we have $\max_{\ell\in\mathcal{S}}[\psi^v_\omega(g(p_{\ell'} \circ p')) - \psi^v_\omega(g_\ell)] > \delta$. Then, we have $\max_{p\in P(s,\Gamma)}[\psi^v_\omega(g(p_{\ell'} \circ p')) - \psi^v_\omega(g(p))] > \delta$ for all $\omega \in \Omega'$ since $\{p_\ell : \ell \in \mathcal{S}\} \subseteq P(s, \Gamma)$. Therefore, we have $\max_{\omega\in\Omega'}\max_{p\in P(s,\Gamma)}[\psi^v_\omega(g(p_{\ell'} \circ p')) - \psi^v_\omega(g(p))] > \delta$, i.e. $\mathrm{MR}(g(p_{\ell'} \circ p'), \mathcal{X}, \Omega') > \delta$. □

Thus, if there exists a label $\ell' \in \mathcal{O}$ such that we have $g_\mathcal{S} \lhd^\delta_\Omega F(\ell')$ at some point of the search procedure, then Proposition 2 ensures that path $p_{\ell'}$ cannot be completed into a solution path with a max regret MR below δ (even if we further restrict the set of feasible weights Ω by asking preference queries to the DM). This result can be used to insert a pruning rule in the search so as to detect faster a solution path with a MR below δ, if it exists. However, for a given set Ω, it may be the case that no such path exists. We introduce now a sufficient condition on $\mathrm{MMR}(g_\mathcal{S}, \Omega)$ to guarantee the existence of a path with MR below δ:

Proposition 3. *If $\mathrm{MMR}(g_\mathcal{S}, \Omega) \leq \delta$ at the end of the search procedure, then $\mathrm{MR}(g(p^*), \mathrm{ND}(\mathcal{X}), \Omega) \leq \delta$, for any solution path $p^* \in \arg\min_{p_\ell : \ell\in\mathcal{S}} \mathrm{MR}(g_\ell, g_\mathcal{S}, \Omega)$.*

Proof. Let $p^* \in \arg\min_{p_\ell:\ell\in\mathcal{S}} \mathrm{MR}(g_\ell, g_\mathcal{S}, \Omega)$. We want to prove that $\mathrm{PMR}(g(p^*), g(p'), \Omega) \leq \delta$ for any solution path p' such that $g(p') \in \mathrm{ND}(\mathcal{X})$. Two cases may occur:

Case 1: There exists $\ell \in \mathcal{S}$ such that $p' = p_\ell$. In that case, we can directly infer the result because $\mathrm{PMR}(g(p^*), g(p'), \Omega) \leq \mathrm{MR}(g(p^*), g_\mathcal{S}, \Omega) = \mathrm{MMR}(g_\mathcal{S}, \Omega) \leq \delta$.

Case 2: There exists no $\ell \in \mathcal{S}$ such that $p' = p_\ell$. In that case, there exists $\mathcal{S}' \subseteq \mathcal{S}$ and a generated label ℓ' such that path $p_{\ell'}$ is a subpath of p' and $g_{\mathcal{S}'} \lhd^\delta_\Omega F(\ell')$. Thus, for all $h \in H(n_{\ell'})$ and all $\omega \in \Omega$, there exists a path $p^h_\omega \in \{p_\ell : \ell \in \mathcal{S}'\}$ such that $\psi^v_\omega(g_{\ell'} + h) - \psi^v_\omega(g(p^h_\omega)) > \delta$. Then, since H is admissible, there exists $h' \in H(n_{\ell'})$ such that $g_{\ell'} + h'$ Pareto-dominates $g(p')$, and since ψ^v_ω is increasing with Pareto-dominance, then we have $\psi^v_\omega(g_{\ell'} + h') \leq \psi^v_\omega(g(p'))$. Therefore, we have $\psi^v_\omega(g(p')) - \psi^v_\omega(g(p^{h'}_\omega)) > \delta$, which can be rewritten $\psi^v_\omega(g(p^{h'}_\omega)) - \psi^v_\omega(g(p')) < -\delta$. Moreover, we necessarily have $\psi^v_\omega(g(p^*)) - \psi^v_\omega(g(p^{h'}_\omega)) \leq \delta$ since $\mathcal{S}' \subseteq \mathcal{S}$ and $\mathrm{MR}(g(p^*), g_\mathcal{S}, \Omega) \leq \delta$. Finally, we have $\psi^v_\omega(g(p^*)) - \psi^v_\omega(g(p')) < 0 \leq \delta$ by summing the two previous inequalities and therefore $\mathrm{PMR}(g(p^*), g(p'), \Omega) \leq \delta$. □

Therefore, if we ensure that $\mathrm{MMR}(g_\mathcal{S}, \Omega)$ is smaller than threshold δ at the end of the procedure, then any solution $p^* \in \arg\min_{p_\ell:\ell\in\mathcal{S}} \mathrm{MR}(g_\ell, g_\mathcal{S}, \Omega)$ satisfies $\mathrm{MR}(g(p^*), \mathrm{ND}(\mathcal{X}), \Omega) \leq \delta$. In order to decrease the minimax regret $\mathrm{MMR}(g_\mathcal{S}, \Omega)$, we can, at any step of the search procedure, ask preference queries to reduce the set of feasible parameters Ω. Indeed, it can easily be checked that $\Omega' \subseteq \Omega$ implies $\mathrm{MMR}(g_\mathcal{S}, \Omega') \leq \mathrm{MMR}(g_\mathcal{S}, \Omega)$. For example, we may use a query selection strategy proposed in [Boutilier *et al.*, 2006] that consits of asking the DM to compare x^* and y^* where x^* is the current MR optimal vector in $g_\mathcal{S}$ and y^* is the worse adversary choice (i.e. the vector maximizing $\mathrm{PMR}(x^*, y, \Omega)$ for $y \in g_\mathcal{S}$). The answer to this query induces a linear constraint that can be used to restrict Ω.

We have implemented this first procedure on state space graphs (numerical results are reported in Section 5). When δ is small, the guarantee on regrets is good but the number of queries is quite important. As δ increases, the number of queries diminishes but the procedure becomes significantly slower due to the size of the uncertainty set Ω that makes the pruning rule less efficient. For this reason, we propose in the next section a sophistication of our procedure using an approximation algorithm, that will be much more efficient while still providing guarantees on MMR values.

4 Combining Approximation and Elicitation

In order to obtain a faster search algorithm, we are going to work on near optimal cost vectors with respect to functions $\psi^v_\omega, \omega \in \Omega$. For this reason we introduce the following dominance relation:

Definition 2 (**$\precsim^\varepsilon_\Omega$-dominance**). $\forall X, Y \subset \mathbb{R}^q_+$, $\forall \varepsilon \geq 0$*:*

$$X \precsim^\varepsilon_\Omega Y \Leftrightarrow \forall y \in Y, \forall \omega \in \Omega, \exists x \in X, (1+\varepsilon)\psi^v_\omega(y) \geq \psi^v_\omega(x)$$

Similarly to $\lhd^\delta_\Omega$- dominance tests, $\precsim^\varepsilon_\Omega$-dominance tests can be performed using linear programming due to the following:

Proposition 4. *For all $X, Y \subset \mathbb{R}^q_+$:*

$$X \precsim^\varepsilon_\Omega Y \;\Leftrightarrow\; \forall y \in Y, \min_{\omega\in\Omega}, \max_{x\in X} \left[(1+\varepsilon)\psi^v_\omega(y) - \psi^v_\omega(x)\right] \geq 0$$

The proof is deliberately omitted because it is very similar to that of Proposition 1. Let us show now that $\precsim^\varepsilon_\Omega$ is a relaxation to the $\lhd^\delta_\Omega$ dominance introduced in Definition 1.

Proposition 5. *For all $X, Y \subset \mathbb{R}^q_+$: $X \lhd^\delta_\Omega Y \Rightarrow X \precsim^\varepsilon_\Omega Y$*

Proof. Consider $X, Y \subseteq \mathbb{R}^q_+$ such that $X \lhd^\delta_\Omega Y$. Let $y \in Y$ and $\omega \in \Omega$. Since $X \lhd^\delta_\Omega Y$, then there exists $x \in X$ such that $\psi^v_\omega(y) - \psi^v_\omega(x) > \delta \geq 0$, i.e. such that $\psi^v_\omega(y) \geq \psi^v_\omega(x)$. Moreover, we have $(1+\varepsilon)\psi^v_\omega(y) \geq \psi^v_\omega(y)$ since $\varepsilon \geq 0$ and $\psi^v_\omega(y) \geq 0$, and therefore $(1+\varepsilon)\psi^v_\omega(y) \geq \psi^v_\omega(x)$. Hence we have $X \precsim^\varepsilon_\Omega Y$. □

As a consequence, using $\precsim^\varepsilon_\Omega$-dominance instead of $\lhd^\delta_\Omega$-dominance to prune open labels in MOA* may reduce the number of generated labels (and probably solution times). However, when using this sharper pruning rule, we loose the guarantee obtained for MR values in Proposition 3. In order to restore a guarantee when using the pruning rule based on the $\precsim^\varepsilon_\Omega$-dominance we propose to work with the following definition of regrets.

Definition 3. *For all $x, y \in \mathcal{X}$:*
$\mathrm{PMR}_\varepsilon(x, y, \Omega) = \max_{\omega\in\Omega}\{(1+\varepsilon)\psi^v_\omega(x) - \psi^v_\omega(y)\}$

$\mathrm{MR}_\varepsilon(x, \mathcal{X}, \Omega) = \max_{y\in\mathcal{X}} \mathrm{PMR}_\varepsilon(x, y, \Omega)$

$\mathrm{MMR}_\varepsilon(\mathcal{X}, \Omega) = \min_{x\in\mathcal{X}} \mathrm{MR}_\varepsilon(x, \mathcal{X}, \Omega)$

These regrets are obviously an extension of the initial notions of regrets introduced in Section 3. When $\varepsilon = 0$, they are identical to the initial definition of regrets. When $\varepsilon > 0$ their definition enables to establish the following counterpart of Proposition 3:

Theorem 1. *If $\mathrm{MMR}_\varepsilon(g_\mathcal{S}, \Omega) \leq (1+\varepsilon)\delta$ at the end of the search procedure, then $\mathrm{MR}(g(p^*), \mathrm{ND}(\mathcal{X}), \Omega) \leq \delta$, for any solution path $p^* \in \arg\min_{p_\ell:\ell\in\mathcal{S}} \mathrm{MR}(g_\ell, g_\mathcal{S}, \Omega)$.*

Proof. Let $p^* \in \arg\min_{p_\ell:\ell\in\mathcal{S}} \mathrm{MR}_\varepsilon(g_\ell, g_\mathcal{S}, \Omega)$. We want to prove that $\mathrm{PMR}(g(p^*), g(p'), \Omega) \leq \delta$ for any solution p' such that $g(p') \in \mathrm{ND}(\mathcal{X})$. Let $\lambda = (1+\varepsilon)\delta$. Two cases may occur:

Case 1: There exists $\ell \in \mathcal{S}$ such that $p' = p_\ell$. In that case, we have $\mathrm{PMR}_\varepsilon(g(p^*), g(p'), \Omega) \leq \lambda$ since $\mathrm{MMR}_\varepsilon(g_\mathcal{S}, \Omega) \leq \lambda$, i.e. $(1+\varepsilon)\psi^v_\omega(g(p^*)) - \psi^v_\omega(g(p')) \leq \lambda$ for all $\omega \in \Omega$. Then, we have $(1+\varepsilon)\psi^v_\omega(g(p^*)) - (1+\varepsilon)\psi^v_\omega(g(p')) \leq \lambda$ since $\varepsilon \geq 0$, and so $\psi^v_\omega(g(p^*)) - \psi^v_\omega(g(p')) \leq \lambda/(1+\varepsilon) = \delta$. Hence we have $\mathrm{PMR}(g(p^*), g(p'), \Omega) \leq \delta$.

Case 2: There exists no $\ell \in \mathcal{S}$ such that $p' = p_\ell$. In that case, there exists $\mathcal{S}' \subseteq \mathcal{S}$ and a generated label ℓ' such that $p_{\ell'}$ is a subpath of p' and $\{g_\ell : \ell \in \mathcal{S}'\} \precsim^\varepsilon_\Omega F(\ell')$. Therefore, for all $\omega \in \Omega$ and all $h \in H(n_{\ell'})$, there exists $p^h_\omega \in \{p_\ell : \ell \in \mathcal{S}'\}$ such that $(1+\varepsilon)\psi^v_\omega(g_{\ell'} + h) \geq \psi^v_\omega(g(p^h_\omega))$. Moreover, since H is admissible, there exists $h' \in H(n_{\ell'})$ such that $g_{\ell'} + h'$ Pareto-dominates $g(p')$, and since ψ^v_ω is compatible with Pareto-dominance, we have $\psi^v_\omega(g_{\ell'} + h') \leq \psi^v_\omega(g(p'))$. Therefore, we have $(1+\varepsilon)\psi^v_\omega(g(p')) \geq \psi^v_\omega(g(p^h_\omega))$. Moreover, since $\mathcal{S}' \subseteq \mathcal{S}$ and $\mathrm{MMR}_\varepsilon(g_\mathcal{S}, \Omega) \leq \lambda$, then we have $(1+\varepsilon)\psi^v_\omega(g(p^*)) - \psi^v_\omega(g(p^h_\omega)) \leq \lambda$. Finally, we obtain $(1+\varepsilon)\psi^v_\omega(g(p^*)) - (1+\varepsilon)\psi^v_\omega(g(p')) \leq \lambda$ from the two previous inequalities, i.e. $\psi^v_\omega(g(p^*)) - \psi^v_\omega(g(p')) \leq \lambda/(1+\varepsilon) = \delta$. Hence, we have $\mathrm{PMR}(g(p^*), g(p'), \Omega) \leq \delta$. □

Therefore, if we obtain $\mathrm{MMR}_\varepsilon(g_\mathcal{S}, \Omega) \leq (1+\varepsilon)\delta$ at the end of the search procedure, then any solution path p^* in $\arg\min_{p_\ell:\ell\in\mathcal{S}} \mathrm{MR}(g_\ell, g_\mathcal{S}, \Omega)$ is such that $\mathrm{MR}(g(p^*), \mathrm{ND}(\mathcal{X}), \Omega) \leq \delta$. In order to decrease $\mathrm{MMR}_\varepsilon(g_\mathcal{S}, \Omega)$, we can, here also, ask preference queries to reduce the set of admissible parameters Ω. This is less straightforward than in Section 3 due to the use of MMR_ε instead of standard MMR. However the query selection strategy used in Section 3 can be adapted as shown by the following proposition:

Proposition 6. *If $\varepsilon \leq \delta/(1-\delta)$, then there exists a questionnaire that enables to reduce Ω in such way that $\mathrm{MMR}_\varepsilon(g_\mathcal{S}, \Omega) \leq (1+\varepsilon)\delta$.*

Proof. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the directed graph defined as follows: $\mathcal{V}$ is a set of nodes, one node denoted v_ℓ per label $\ell \in \mathcal{S}$, and $\mathcal{E}$ is a set of arcs where node v_ℓ is linked to node $v_{\ell'}$ if and only if $\psi^v_\omega(g_\ell) \leq \psi^v_\omega(g_{\ell'})$ for all $\omega \in \Omega$. Note that for every cycle of type $\langle v_{\ell_1}, \ldots, v_{\ell_k}, v_{\ell_1}\rangle$ in $\mathcal{G}$, we necessarily have $\psi^v_\omega(g_{\ell_1}) = \ldots = \psi^v_\omega(g_{\ell_k})$ for all $\omega \in \Omega$. Hence, each maximal cycle can be reduced to a single node. Let $\widehat{\mathcal{G}} = (\widehat{\mathcal{V}}, \widehat{\mathcal{E}})$ be the directed acyclic graph obtained after these reductions. Then, consider the following elicitation procedure.

While graph $\widehat{\mathcal{G}}$ have more than one source:

- Determine x^*_0 as one solution minimizing $\mathrm{MR}(x, g_\mathcal{S}, \Omega)$ and y^*_0 a solution that maximizes $\mathrm{PMR}(x^*_0, y, \Omega)$. Let x^* and y^* be respectively an ancestor of x^*_0 and y^*_0 that have no predecessor. If x^*_0 has no ancestor, then $x^* = x^*_0$ (the same applies to y^*_0).
- Ask the DM to compare the two solutions associated with x^* and y^*.
- Update graph $\mathcal{G}$ by inserting new arcs induced by the new preference information obtained and $\widehat{\mathcal{G}}$ accordingly.

At the end, we obtain a connected digraph, the source of which is denoted v_{ℓ^*}. By construction, for all $\ell \in \mathcal{S}$, we have $\psi^v_\omega(g_{\ell^*}) \leq \psi^v_\omega(g_\ell)$ for all $\omega \in \Omega$. Therefore, we have $\mathrm{PMR}_\varepsilon(g_{\ell^*}, g_\ell, \Omega) = \max_{\omega\in\Omega}\{(1+\varepsilon)\psi^v_\omega(g_{\ell^*}) - \psi^v_\omega(g_\ell)\} \leq \max_{\omega\in\Omega}\{(1+\varepsilon)\psi^v_\omega(g_{\ell^*}) - \psi^v_\omega(g_{\ell^*})\} = \varepsilon \times \max_{\omega\in\Omega} \psi^v_\omega(g_{\ell^*})$. Then, since $\psi^v_\omega(g_{\ell^*}) \leq 1$ for all $\omega \in \Omega$, we have $\mathrm{PMR}_\varepsilon(g_{\ell^*}, g_\ell, \Omega) \leq \varepsilon$. Hence $\mathrm{PMR}_\varepsilon(g_{\ell^*}, g_\ell, \Omega) \leq (1+\varepsilon)\delta$ directly follows from $\varepsilon \leq \delta/(1-\delta)$. Thus $\mathrm{MR}_\varepsilon(g_{\ell^*}, g_\mathcal{S}, \Omega) \leq (1+\varepsilon)\delta$ and $\mathrm{MMR}_\varepsilon(g_\mathcal{S}, \Omega) \leq (1+\varepsilon)\delta$. □

5 Numerical Tests

We have evaluated the performance of the two algorithms respectively presented in Section 3 and 4 in terms of computation times (in seconds) and number of queries. Results are obtained by averaging over 30 runs and linear optimizations are performed using the Gurobi library of Java. The algorithm based on MR minimization is denoted R^* hereafter whereas the one based on MR_ε minimization is denoted R^*_ε.

In a series of experiments, we consider instances of graphs $G = (N, A)$ where all nodes in N are uniformly drawn in the two dimension grid $\{1, \ldots, 1000\} \times \{1, \ldots, 1000\}$, but source node s and goal node γ are respectively located in $(1, 500)$ and $(1000, 500)$. Each node is linked to 30 randomly chosen nodes and the associated cost vectors are randomly drawn using Gaussian distributions parametrized according to Euclidean distances. For each node $n \in N$, we set

$H(n) = \{I(n)\}$ where $I(n) = (I_1(n), \ldots, I_q(n))$ is the ideal point defined by $I_i(n) = \min_{p \in P(n,\gamma)} g_i(p)$ for all $i \in Q$. We consider S-shaped disutility functions v_i, $i \in Q$, of the form:

$$v_i(x_i) = \frac{1}{1 + \mathrm{e}^{-a_i(x_i - b_i)}}$$

where x_i is the i^{th} component of cost vector x; a_i and b_i are parameters enabling respectively to control the amplitude of the 'S' and the position of the 'S' along the i-th criterion.

To evaluate the impact of the model complexity (in terms of number of parameters), we consider additive utilities (U_ω^v) and Choquet integrals (C_ω^v) of type:

$$C_\omega^v(x) = \sum_{i \in Q} m_i v_i(x_i) + \sum_{i,j \in Q: i<j} m_{i,j} \min\{v_i(x_i), v_j(x_j)\}$$

This is indeed a Choquet expected utility associated with the capacity $\omega(X) = \sum_{i \in X} m_i + \sum_{i,j \in X: i<j} m_{i,j}$ for $X \subseteq Q$. This specific subclass of capacity is said to be 2-additive because it only involves a number of parameters which is quadratic in the number of criteria.

For algorithm R_ε^*, we have estimated the value of ε that enables a balanced trade-off between the computation time and the number of queries. Table 1 shows that the larger parameter ε, the smaller the number of queries and the computation time. Hence, the best option is definitely to set ε to its maximum feasible value, i.e. $\varepsilon = \delta/(1-\delta)$. Therefore, in the following experiments, we only consider this value for ε.

ε	0	$\delta/3(1-\delta)$	$2\delta/3(1-\delta)$	$\delta/(1-\delta)$
time	6.50	5.66	4.99	4.38
queries	13.80	12.53	11.63	10.43

Table 1: Performance of R_ε^* ($\delta = 0.2, q = 10, |N| = 1000, C_\omega^v$).

Then, we have compared R^* and R_ε^* algorithms. In Table 2, we can see that computation times drastically increase with respect to δ for the R^* algorithm. In other words, when lowering δ (the required guarantee of quality), algorithm R^* reduces the number of queries but increases significantly computation times, a drawback that does not appear for R_ε^*. Moreover, R_ε^* is much faster than R^*, e.g. 3000 times faster for $q = 2$ and $\delta = 0.1$. The time difference between the two algorithms increases with q, the number of criteria.

Finally, we have investigated the impact of the number of parameters, on the performance of R_ε^*. As it could be expected, the number of queries and the computation times needed for elicitating the 2-additive Choquet model are more important than for elicitating the additive utility model (the former model involves $q(q+1)/2$ parameters, the latter only q). This is the price to pay for higher descriptive and prescriptive possibilities but in any case the overall number of queries remains quite admissible as the number of criteria increases.

		$\delta = 0.01$		$\delta = 0.05$		$\delta = 0.1$	
q	method	time	queries	time	queries	time	queries
2	R^*	0.6	3.0	7.9	1.9	309.7	1.0
2	R_ε^*	0.4	3.2	0.2	2.0	0.1	1.4
4	R^*	3.6	12.1	68.7	7.4	1537.0	4.4
4	R_ε^*	1.9	12.5	1.0	8.1	0.6	5.1

Table 2: Comparison of R^* and R_ε^* ($|N| = 500$, C_ω^v).

			$\delta = 0.05$		$\delta = 0.1$	
q	model	$\|N\|$	time	queries	time	queries
2	U_ω^v, C_ω^v	500	0.2	1.8	0.1	1.1
2	U_ω^v, C_ω^v	1000	0.2	2.0	0.1	1.4
5	U_ω^v	500	1.1	8.1	0.7	5.6
5	C_ω^v	500	2.9	13.5	1.5	8.9
5	U_ω^v	1000	1.4	8.3	0.9	5.6
5	C_ω^v	1000	3.7	13.4	2.0	8.1
10	U_ω^v	500	7.8	22.3	3.4	13.7
10	C_ω^v	500	66.9	47.5	18.1	27.2
10	U_ω^v	1000	10.0	26.9	4.3	16.0
10	C_ω^v	1000	138.2	53.3	43.0	30.5

Table 3: Elicitating U_ω^v versus elicitating C_ω^v.

6 Conclusion

We have introduced a new approach combining near-admissible state-space search and incremental elicitation procedures to solve search problems in vector-valued graphs. It makes it possible to elicit the weighting parameters of non-linear models such as the additive utility model and the Choquet expected utility. Our approach is based on a sophistication of MOA* search and U^* search involving new pruning rules implemented with an LP solver. The search procedure is interweaved with an incremental elicitation procedure allowing to approximate, more and more accurately, the preference parameters controlling the importance of criteria, or sets of criteria, and thus the value system of the decision maker.

The numerical tests reported show the efficiency of this approach both in terms of number of queries and in terms of solution times. The elicitation procedure based on MR_ε regrets minimization is shown to provide robust solutions, i.e. solutions with a gap to optimality guaranteed to fall below a desired threshold δ.

A natural extension of this work would be to integrate the elicitation of one dimensional utility functions in the whole regret minimization process, instead of eliciting them in a preliminary stage. This seems to be a challenging question because in this case, regret minimization will require to solve quadratic optimization problems at every step of the search procedure.

Acknowledgments

This work is part of the ELICIT project supported by the French National Research Agency through the Idex Sorbonne Universités under grant ANR-11-IDEX-0004-02.

References

[Bana e Costa and Vansnick, 2000] C.A. Bana e Costa and J.C. Vansnick. Cardinal value measurement with MACBETH. In S.H. Zanakis, G. Doukidis, and C. Zopounidis, editors, *Decision Making: Recent Developments and Worldwide Applications*, pages 317–329. Kluwer, 2000.

[Benabbou and Perny, 2015] Nawal Benabbou and Patrice Perny. Incremental Weight Elicitation for Multiobjective State Space Search. In *Proceeding of AAAI-15*, 2015.

[Benabbou *et al.*, 2014] Nawal Benabbou, Patrice Perny, and Paolo Viappiani. Incremental Elicitation of Choquet Capacities for Multicriteria Decision Making. In *Proceedings of ECAI-14*, pages 87–92, 2014.

[Boutilier *et al.*, 2006] Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Constraint-based Optimization and Utility Elicitation using the Minimax Decision Criterion. *Artifical Intelligence*, 170(8–9):686–713, 2006.

[Boutilier, 2002] Craig Boutilier. A POMDP Formulation of Preference Elicitation Problems. pages 239–246, 2002.

[Dasgupta *et al.*, 1995] P. Dasgupta, P.P. Chakrabarti, and S.C. DeSarkar. Utility of *pathmax* in partial order heuristic search. *J. of algorithms*, 55:317–322, 1995.

[Dasgupta *et al.*, 1996a] Pallab Dasgupta, PP Chakrabarti, and SC DeSarkar. Multiobjective heuristic search in and/or graphs. *Journal of algorithms*, 20(2):282–311, 1996.

[Dasgupta *et al.*, 1996b] Pallab Dasgupta, PP Chakrabarti, and SC DeSarkar. Searching game trees under a partial order. *Artificial Intelligence*, 82(1-2):237–257, 1996.

[Dery *et al.*, 2014] Lihi Naamani Dery, Meir Kalech, Lior Rokach, and Bracha Shapira. Reaching a joint decision with minimal elicitation of voter preferences. *Information Sciences*, 278:466–487, 2014.

[Ding and Lin, 2013] Ning Ding and Fangzhen Lin. Voting with partial information: what questions to ask? In *Proceedings of AAMAS-13*, pages 1237–1238, 2013.

[Fishburn, 1967] Peter C Fishburn. Methods of estimating additive utilities. *Management science*, 13(7):435–453, 1967.

[Fishburn, 1968] Peter C Fishburn. Utility theory. *Management science*, 14(5):335–378, 1968.

[Grabisch *et al.*, 2009] M. Grabisch, J.-L. Marichal, R. Mesiar, and E. Pap. *Aggregation Functions*. Cambridge University Press, New-York, 2009.

[Hart *et al.*, 1968] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. and Cyb.*, SSC-4 (2):100–107, 1968.

[Kalech *et al.*, 2010] Meir Kalech, Sarit Kraus, and Gal A. Kaminka. Practical voting rules with partial information. *Autonomous Agents and Multi-Agent Systems*, 22(1):151–182, 2010.

[Keeney and Raiffa, 1976] R.L. Keeney and H. Raiffa. *Decisions with multiple objectives: Preferences and value tradeoffs*. J. Wiley, New York, 1976.

[Konczak and Lang, 2005] Kathrin Konczak and Jerome Lang. Voting procedures with incomplete preferences. In *Proceedings of IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling*, volume 20, 2005.

[Krantz *et al.*, 1971] D.H. Krantz, R.D. Luce, P. Suppes, and A. Tversky. *Foundations of measurement*. Academic Press, New York, 1971.

[Lang *et al.*, 2012] Jérôme Lang, Maria Silvia Pini, Francesca Rossi, Domenico Salvagnin, Kristen Brent Venable, and Toby Walsh. Winner determination in voting trees with incomplete preferences and weighted votes. *Autonomous Agents and Multi-Agent Systems*, 25(1):130–157, 2012.

[Lu and Boutilier, 2011] Tyler Lu and Craig Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of IJCAI*, pages 287–293, 2011.

[Lu and Boutilier, 2013] Tyler Lu and Craig Boutilier. Multi-winner social choice with incomplete preferences. In *Proceedings of IJCAI*, pages 263–270, 2013.

[Mandow and de la Cruz, 2005] L. Mandow and J.L. Prez de la Cruz. A new approach to multiobjective A^* search. In *Proceedings of IJCAI-05*, pages 218–223, 2005.

[Marinescu *et al.*, 2013] Radu Marinescu, Abdul Razak, and Nic Wilson. Multi-objective constraint optimization with tradeoffs. In *CP 2013*, pages 497–512, 2013.

[Marinescu, 2010] Radu Marinescu. Best-first vs. depth-first AND/OR search for multi-objective constraint optimization. In *ICTAI 2010*, pages 439–446, 2010.

[Perny and Spanjaard, 2003] Patrice Perny and Olivier Spanjaard. On preference-based Search in State Space Graphs. In *Proc. of AAAI-2002 Conference*, pages 751–756, 2003.

[White III *et al.*, 1984] C. C. White III, A. P. Sage, and S. Dozono. A model of multiattribute decisionmaking and trade-off weight determination under uncertainty. *IEEE Transactions on Systems, Man, and Cybernetics*, 14(2):223–229, 1984.

[Savage, 1954] Leonard J. Savage. *The Foundations of Statistics*. Wiley, 1954.

[Schmeidler, 1986] D. Schmeidler. Integral representation without additivity. *Proceedings of the American Mathematical Society*, 97(2):255–261, 1986.

[Stewart and White III, 1991] B.S. Stewart and C.C. White III. Multiobjective A^*. *Journal of the Association for Computing Machinery*, 38(4):775–814, 1991.

[Wang and Boutilier, 2003] T. Wang and C. Boutilier. Incremental Utility Elicitation with the Minimax Regret Decision Criterion. In *Proc. of IJCAI-03*, pages 309–316, 2003.

[Xia and Conitzer, 2011] Lirong Xia and Vincent Conitzer. Determining possible and necessary winners given partial orders. *Journal of Artificial Intelligence Research (JAIR)*, 41:25–67, 2011.