24th European Conference on Artificial Intelligence - ECAI 2020
Santiago de Compostela, Spain

Minimality and Comparison of Sets of Multi-Attribute Vectors

Federico Toffano and **Nic Wilson** [1]

Abstract. In a decision-making problem, there is often some uncertainty regarding the user preferences. We assume a parameterised utility model, where in each scenario we have a utility function over alternatives, and where each scenario represents a possible user preference model consistent with the input preference information. With a set A of alternatives available to the decision maker, we can consider the associated utility function, expressing, for each scenario, the maximum utility among the alternatives. We consider two main problems: firstly, finding a minimal subset of A that is equivalent to it, i.e., that has the same utility function. We show that for important classes of preference models, the set of so-called possibly strictly optimal alternatives is the unique minimal equivalent subset. Secondly, we consider how to compare A to another set of alternatives B, where A and B correspond to different initial decision choices. We derive mathematical results that allow different computational techniques for these two problems, using linear programming, and especially, with a novel approach using the extreme points of the epigraph of the utility function.

1 INTRODUCTION

In a decision-making problem, there can be uncertainty regarding the user preferences. Suppose that, in a particular situation, A is the set of alternatives that are available to the decision maker. This is interpreted in a disjunctive fashion, in that the user is free to choose any element α of A. However, as is common, we do not know precisely the user's preferences. The preference information available to the system is represented in terms of a set of user preference models, parameterised by a set (of *scenarios*) $\mathcal{W}$ where, associated with each scenario $w \in \mathcal{W}$, is a (real-valued) utility function f_w over alternatives.

Each element w of $\mathcal{W}$ is viewed as a possible model of the user's preferences that is consistent with the preference information we know. If we knew that w were the true scenario, so that f_w represents the user's preferences over alternatives, then we would be able to choose a best element of A with respect to f_w leading to a utility value $Ut_A(w) = \max_{\alpha \in A} f_w(\alpha)$. However, the situation can be ambiguous given a non-singleton set $\mathcal{W}$ of possible user models or scenarios.

The set $\mathcal{W}$ incorporates what we know about the user preferences; for example, if we have learned that the user regards alternative β as at least as good as alternative γ, then $\mathcal{W}$ will only include scenarios w such that $f_w(\beta) \geq f_w(\gamma)$.

[1] Insight Centre for Data Analytics, School of Computer Science and Information Technology, University College Cork, Cork, Ireland {federico.toffano,nic.wilson}@insight-centre.org

The utility function f_w may be based on a decomposition of utility, using, for example, an additive representation for a combinatorial problem (e.g., [23, 32, 25]). Also, $f_w(\alpha)$ could represent the expected utility of alternative α given that w is the correct user model, based on a probabilistic model with parameter w, for example in a multi-objective influence diagram [15, 24, 26], with α corresponding to a policy.

We consider, in particular, the following related pair of questions:

(1) Are there elements of A that can be eliminated unproblematically? In particular, is there a strict subset A' of A that is equivalent to A?

(2) Given a choice between one situation, in which the available alternatives are A, and another situation, in which alternatives B are available, is A at least as good as B in every scenario?

Regarding (1), we need to be able to eliminate unimportant choices, to make the list of options manageable, in particular, if we want to display the alternatives to the user. We interpret this as finding a minimal subset A' of A such that $Ut_A(w) = Ut_{A'}(w)$ for every scenario $w \in \mathcal{W}$.

Question (2) concerns a case in which the user may have a choice between (I) being able to obtain any of the set of alternatives A, and (II) any alternative in B (and thus, the user could obtain any alternative in $A \cup B$). Sets A and B may correspond to different choices $X = a$ and $X = b$ of a fundamental variable X, and determining that A dominates B may lead us to exclude $X = b$, thus simplifying the problem. For instance, A might correspond to hotels in Paris, and B to hotels in Lisbon, for a potential weekend away. We want to be able to determine if one of these clearly dominates the other; if, for instance, A dominates B, then there may be no need for the system and the user to further consider B, and, for example, may focus on Paris rather than Lisbon. We interpret this task as determining if in every scenario the utility A is at least that for B, i.e., $Ut_A(w) \geq Ut_B(w)$.

The focus of this paper is to determine important properties of the dominance and equivalence relations, and to derive computational procedures, in order to find a minimal equivalent subset, and testing dominance between A and B; we also determine properties and a computational technique for a form of maximum regret, that can be viewed as a degree of dominance, and which corresponds to *setwise max regret* defined in [32]. The main computational procedures are based on linear programming (LP), or, alternatively, a novel method using the extreme points of the epigraph of the utility function (which we abbreviate to EEU).

From the computational perspective we focus especially on the case in which alternatives are represented as multi-attribute utility vectors, based on a weighted average user preference model. Each utility vector is then an element of $I\!R^p$, representing a number p of scales of utility (or objectives); each scenario w is a normalised non-negative vector, with $\alpha \succcurlyeq_w \beta$ if and only if the weighted sum of α with respect to w is at least that of β. An input preference of α over β then leads to a linear constraint on the weights vector w, and we can define the set of consistent preference models $\mathcal{W}$ as the convex polytope generated by a set of input preferences of this form.

We show that for important classes of preference models, the set of Possibly Strictly Optimal (PSO) alternatives is the unique minimal equivalent subset. Furthermore, the PSO operator can be used to filter query sets to avoid the potential of an inconsistent answer.

Section 2 gives the formal setup, defining dominance relations, and giving basic properties. Section 3 discusses related work. Section 4 considers the problem of reducing the size of a set A, whilst maintaining equivalence. Section 5 defines a form of maximum regret in this context, shows how it relates to dominance, and gives properties that will be useful for computation. Section 6 discusses the importance of the possibly optimal and possibly strictly optimal alternatives in incremental preference elicitation. Section 7 describes the EEU method. Section 8 brings together the computational techniques for the weighted multi-attribute utility case. Sections 9 and 10 describe the implementation and experimental testing, and Section 11 concludes.

The online longer version [30] includes proofs, using auxiliary lemmas, and further details about the experiments.

2 TERMINOLOGY AND BASIC PROPERTIES

We consider a (possibly infinite) set Ω of alternatives, and another set $\mathcal{U}$, the elements of which we call *scenarios*, that corresponds with the set of user preference models. With each scenario $w \in \mathcal{U}$ is associated a utility function f_w on Ω, i.e., a function from Ω to $\mathbb{R}$; this gives rise to a total pre-order $\succcurlyeq_w$ on Ω given by $\alpha \succcurlyeq_w \beta \iff f_w(\alpha) \geq f_w(\beta)$, for $\alpha, \beta \in \Omega$.

We do not *a priori* assume anything about the functions f_w; however certain mathematical results make additional assumptions, such as continuity with respect to w. Of particular interest in this paper is the case when $f_w(\alpha)$ is a linear function of w, where $\mathcal{U} = \mathbb{R}^p$ for some p, and so $f_w(\alpha)$ can be written as $\sum_{i=1}^{p} \alpha_i w(i) = (\alpha_1, \ldots, \alpha_p) \cdot w$, for some reals α_i, with α_i representing how good alternative α is with respect to objective/criterion i. Then α can be identified with the vector $(\alpha_1, \ldots, \alpha_p) \in \mathbb{R}^p$.

For $\mathcal{W} \subseteq \mathcal{U}$ we define relation $\succcurlyeq_\mathcal{W}$ on Ω by $\alpha \succcurlyeq_\mathcal{W} \beta \iff$ for all $w \in \mathcal{W}$, $\alpha \succcurlyeq_w \beta$. Thus, $\alpha \succcurlyeq_\mathcal{W} \beta$ if and only if α is at least as good as β in every scenario in $\mathcal{W}$. We define $\succ_\mathcal{W}$ to be the strict part of $\succcurlyeq_\mathcal{W}$, i.e., for $\alpha, \beta \in \Omega$, $\alpha \succ_\mathcal{W} \beta$ if and only if $\alpha \succcurlyeq_\mathcal{W} \beta$ and $\beta \not\succcurlyeq_\mathcal{W} \alpha$. Relation $\succ_\mathcal{W}$ is transitive and acyclic. We define equivalence relation $\equiv_\mathcal{W}$ to be the symmetric part of $\succcurlyeq_\mathcal{W}$, given by $\alpha \equiv_\mathcal{W} \beta$ if and only if $\alpha \succcurlyeq_\mathcal{W} \beta$ and $\beta \succcurlyeq_\mathcal{W} \alpha$.

Let $\mathcal{M}$ be the set of finite subsets of Ω.

The utility function associated with $A \in \mathcal{M}$: we define, for $w \in \mathcal{U}$, $Ut_A(w)$ to be $\max_{\alpha \in A} f_w(\alpha)$.

Dominance relation between sets: For subset $\mathcal{W}$ of $\mathcal{U}$, we define binary relation $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$ on $\mathcal{M}$ as follows. Consider any $A, B \in \mathcal{M}$.

- $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ if and only if for all $w \in \mathcal{W}$ and for all $\beta \in B$ there exists $\alpha \in A$ such that $\alpha \succcurlyeq_w \beta$.

It is easy to see that $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ if and only if for all $w \in \mathcal{W}$, $Ut_A(w) \geq Ut_B(w)$. Relation $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$ is transitive and satisfies obvious monotonicity properties with respect to A, B and $\mathcal{W}$. We also have $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ if and only if $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} \{\beta\}$ holds for each $\beta \in B$. Define equivalence relation $\equiv^{\mathcal{W}}_{\forall\forall\exists}$ by $A \equiv^{\mathcal{W}}_{\forall\forall\exists} B$ if and only if for all $w \in \mathcal{W}$, $Ut_A(w) = Ut_B(w)$. Thus, $\equiv^{\mathcal{W}}_{\forall\forall\exists}$ is the symmetric part of $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$, with $A \equiv^{\mathcal{W}}_{\forall\forall\exists} B$ if and only if $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ and $B \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} A$.

One can also consider a (strong form of) strict dominance $A \gg^{\mathcal{W}}_{\forall\forall\exists} B$ defined as for all $w \in \mathcal{W}$, $Ut_A(w) > Ut_B(w)$; this corresponds with the dominance relation defined in Definition 2 of [2].

Example. Let $A = \{(11,1), (10,4), (7,5), (6,6), (4,7)\}$ and $B = \{(11,2), (8,5)\}$ be sets of utility vectors of hotels in Paris and Lisbon respectively. For example, the first value of each utility vector could be a score for the location and the second value could be a score for cleanliness, the higher the score, the better. We assume linear utility functions with $f_w(\alpha) = w \cdot \alpha$. Let $\mathcal{U} = \{(w_1, w_2) : w_1, w_2 \geq 0 \,\&\, w_1 + w_2 = 1\}$, representing different normalised weightings of the two criteria. We assume that the user has an associated weights vector that is unknown and we want to recommend to the user a trip to Paris or Lisbon based on her preferences on the available hotels. Suppose then that we ask to the user her preference between the hotel with utility $(10,4)$ and the hotel with utility $(11,2)$. An input preference of $(10,4)$ over $(11,2)$ implies $w \cdot (10,4) \geq w \cdot (11,2)$ and so $2w_2 \geq w_1$ and thus, $w_1 \leq \frac{2}{3}$, leading to the set of scenarios $\mathcal{W} = \{(w_1, w_2) : w_1 + w_2 = 1 \,\&\, 0 \leq w_1 \leq \frac{2}{3}\}$. This example is illustrated in Figure 1, and it is easy to see that $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ since for $0 \leq w_1 \leq \frac{1}{3}$ there is no line above the line associated to $(4,7) \in A$, and for $\frac{1}{3} \leq w_1 \leq \frac{2}{3}$ there is no line above the line associated to $(10,4) \in A$, i.e., $\nexists \beta \in B$ s.t. $f_w(\beta) > Ut_A(w)$ for any $w \in \mathcal{W}$. Therefore in this case we can recommend to the user the trip to Paris.

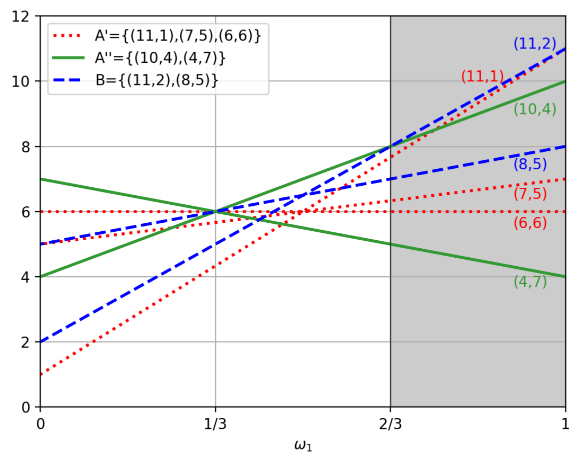

Figure 1. $f_w(\alpha)$ and $f_w(\beta)$ for each $\alpha \in A$ and $\beta \in B$, where $A = A' \cup A'' = \{(11,1), (10,4), (7,5), (6,6), (4,7)\}$, $B = \{(11,2), (8,5)\}$ and $w \in \mathcal{W} = \{(w_1, w_2) : w_1 + w_2 = 1 \,\&\, 0 \leq w_1 \leq \frac{2}{3}\}$.

3 RELATED WORK

Multiattribute utility theory (MAUT) [23] involves numerical representations of user preferences with respect to alternatives evaluated over multiattribute spaces. Imprecisely specified multiattribute utility theory (ISMAUT) [36] is one of the earliest attempt to deal with parameterised utility information representing user preferences with linear inequalities and reducing the set of alternatives to those that are not dominated by any other alternative. Related research such as [20] and [35], deals with similar issues.

A major division in recent work on parameterised user preference models is whether a Bayesian model is assumed over the scenarios (corresponding to the different user preference models), or if there is a purely qualitative (logical) representation of the uncertainty over scenarios, where all we represent is that the scenario is in a set $\mathcal{W}$. Bayesian approaches include [14, 9, 34, 8]. Work involving a qualitative uncertainty representation includes [10, 32, 12, 25, 6]. Linear imprecise preference models, including those based on a simple form of MAUT model, have been considered in work such as [12, 28, 24, 22] including in a conversational recommender system context [13, 32, 33].

In Section 4 we consider a number of operators representing different notions of optimality. The set $\mathrm{UD}_{\mathcal{W}}(A)$, a natural generalisation of the Pareto-optimal elements, appears in many contexts, e.g., [25, 22]. Possibly optimal (also known as *potentially optimal*) elements have been considered in many publications, such as [20, 1, 18, 19, 40, 3, 2, 5]. The Possibly Strictly Optimal set $\mathrm{PSO}_{\mathcal{W}}(A)$ has been considered much less [39, 27, 38]. Regret-based decision making has a long history, with recent work in AI including [10, 12, 6]. We describe in Section 5 the relationship between the dominance relation $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$ and setwise max regret [32].

4 FILTERING A AND MINIMAL EQUIVALENT SUBSETS

In this section we consider the question, raised earlier, regarding replacing A with an equivalent subset of A, i.e., filtering out elements of A that are redundant.

We consider different natural operators for filtering (forms of which have been studied in the literature), namely, $\mathrm{UD}_{\mathcal{W}}(A)$, which removes strictly dominated alternatives from A, and $\mathrm{PO}_{\mathcal{W}}(A)$, which removes alternatives that are not possibly optimal, i.e., not optimal with respect to any scenario in $\mathcal{W}$, and a refined variation, $\mathrm{PSO}_{\mathcal{W}}(A)$.

Setwise-minimal equivalent subsets: We may want to reduce A to an equivalent subset that cannot be reduced any further. We define $\mathrm{SME}_{\mathcal{W}}(A)$ to be the set of subsets B of A that are setwise-minimal equivalent to A, i.e., such that $B \equiv^{\mathcal{W}}_{\forall\forall\exists} A$ and there does not exist any strict subset C of B such that $C \equiv^{\mathcal{W}}_{\forall\forall\exists} A$. Theorem 1 determines when there is a unique such subset. In Section 4.2 we give a simple method for determining setwise-minimal equivalent subsets.

Equivalence-free: we say that A ($\in \mathcal{M}$) is $\equiv_{\mathcal{W}}$-free (or *equivalence-free*) if for all $\alpha, \beta \in A$, we have $\alpha \not\equiv_{\mathcal{W}} \beta$. One can reduce any A to an equivalence-free set A' by including exactly one element in A' of each $\equiv_{\mathcal{W}}$-equivalence class in A.

The Undominated Operator $\mathrm{UD}_{\mathcal{W}}$**:** For $A \in \mathcal{M}$ we define $\mathrm{UD}_{\mathcal{W}}(A)$ to be the set of $\alpha \in A$ such that there does not exist $\gamma \in A$ such that $\gamma \succ_{\mathcal{W}} \alpha$. Thus, the element α of A is not in $\mathrm{UD}_{\mathcal{W}}(A)$ if and only if there exists some $\gamma \in A$ such that γ is at least as good as α in every scenario, and strictly better in at least one scenario. The set $\mathrm{UD}_{\mathcal{W}}(A)$ is a natural generalisation of the Pareto-optimal elements, and is sometimes referred to as the set of *undominated* elements in A.

4.1 The Operators $\mathrm{PO}_{\mathcal{W}}$ and $\mathrm{PSO}_{\mathcal{W}}$

Possibly Optimal Set $\mathrm{PO}_{\mathcal{W}}(A)$**:** for each $w \in \mathcal{U}$ and $A \in \mathcal{M}$ we define $\mathrm{O}_w(A)$ to be all elements α of A that are optimal in A in scenario w, i.e., such that for all $\beta \in A$, $\alpha \succcurlyeq_w \beta$. For $\mathcal{W} \subseteq \mathcal{U}$ we define $\mathrm{PO}_{\mathcal{W}}(A)$ to be $\bigcup_{w \in \mathcal{W}} \mathrm{O}_w(A)$, the set of alternatives that are optimal in some scenario, i.e., optimal for some consistent user preference model.

Possibly Strictly Optimal Set $\mathrm{PSO}_{\mathcal{W}}(A)$**:** we define $\mathrm{SO}^{\mathcal{W}}_w(A)$ to be all elements α of A such that $\alpha \succ_w \beta$, for all $\beta \in A$ with $\beta \not\equiv_{\mathcal{W}} \alpha$. These elements α are said to be *strictly optimal in scenario* w. We then define $\mathrm{PSO}_{\mathcal{W}}(A)$, the set of *possibly strictly optimal* elements, to be $\bigcup_{w \in \mathcal{W}} \mathrm{SO}^{\mathcal{W}}_w(A)$, i.e., all the elements that are strictly optimal in some scenario in $\mathcal{W}$. For equivalence-free A, $\mathrm{PSO}_{\mathcal{W}}(A)$ consists of all alternatives $\alpha \in A$ which are uniquely optimal in some scenario $w \in \mathcal{W}$ (i.e., $\mathrm{O}_w(A) = \{\alpha\}$). It can be easily seen that $\mathrm{PSO}_{\mathcal{W}}(A) \subseteq \mathrm{PO}_{\mathcal{W}}(A) \cap \mathrm{UD}_{\mathcal{W}}(A)$.

Definition of $\mathrm{Opt}^A_{\mathcal{W}}(\alpha)$**:** We define, for $\alpha \in A$, $\mathrm{Opt}^A_{\mathcal{W}}(\alpha)$ to consist of all scenarios $w \in \mathcal{W}$ in which α is optimal, i.e., $\alpha \in \mathrm{O}_w(A)$. Thus, $\alpha \in \mathrm{PO}_{\mathcal{W}}(A) \iff \mathrm{Opt}^A_{\mathcal{W}}(\alpha) \neq \emptyset$. It can be seen that for $B \subseteq A$, we have $B \equiv^{\mathcal{W}}_{\forall\forall\exists} A$ if and only if $\bigcup_{\beta \in B} \mathrm{Opt}^A_{\mathcal{W}}(\beta) = \mathcal{W}$.

Example continued: We have the set of undominated elements $\mathrm{UD}_{\mathcal{W}}(A) = \{(10,4),(7,5),(6,6),(4,7)\}$. Abbreviating w to just its first component w_1 we have $\mathcal{W} = [0, \frac{2}{3}]$, and $\mathrm{Opt}^A_{\mathcal{W}}(10,4) = [\frac{1}{3}, \frac{2}{3}]$; $\mathrm{Opt}^A_{\mathcal{W}}(6,6) = \{\frac{1}{3}\}$, $\mathrm{Opt}^A_{\mathcal{W}}(4,7) = [0, \frac{1}{3}]$ and $\mathrm{Opt}^A_{\mathcal{W}}(11,1) = \mathrm{Opt}^A_{\mathcal{W}}(7,5) = \emptyset$. Thus, $\mathrm{PO}_{\mathcal{W}}(A) = \{(10,4),(6,6),(4,7)\}$. We have $\mathrm{PSO}_{\mathcal{W}}(A) = \{(10,4),(4,7)\}$ and $\{(10,4),(4,7)\} \equiv^{\mathcal{W}}_{\forall\forall\exists} A$. The $\mathrm{PSO}_{\mathcal{W}}$ operator thus leads here to stronger filtering than the $\mathrm{PO}_{\mathcal{W}}$ operator. In Figure 1 we can see a graphical interpretation of $\mathrm{Opt}^A_{\mathcal{W}}((4,7)) = [0, \frac{1}{3}]$, i.e., $w_1 \in [0, \frac{1}{3}]$ is an interval in which there is no line strictly above the line associated to $(4,7)$. We have $(4,7) \in \mathrm{PSO}_{\mathcal{W}}(A)$ because for any $w_1 \in [0, \frac{1}{3})$ the line associated to $(4,7)$ is (strictly) above all the other lines, and $(6,6) \in \mathrm{PO}_{\mathcal{W}}(A)$ because at $w = \frac{1}{3}$ there is no line (strictly) above the line associated to $(6,6)$. □

Theorem 1 below gives some relationships between PSO, SME and the dominance relation $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$, for equivalence-free A. Any setwise-minimal equivalent subset of A contains $\mathrm{PSO}_{\mathcal{W}}(A)$, the set of possibly strictly optimal elements. The latter set is equivalent to A if and only if there is a unique minimal equivalent subset, which is thus equal to $\mathrm{PSO}_{\mathcal{W}}(A)$.

The condition that $\mathrm{PSO}_{\mathcal{W}}(A)$ is equivalent to A holds in the linear multi-objective case considered in Section 8 below (see Theorem 2), and so then $\mathrm{PSO}_{\mathcal{W}}(A)$ is the unique minimal equivalent subset of A. Part (ii) implies that the relation $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$ can be used for computing $\mathrm{PSO}_{\mathcal{W}}(A)$.

Theorem 1 *Assume that A ($\in \mathcal{M}$) is $\equiv_{\mathcal{W}}$-free and let $\mathcal{W} \subseteq \mathcal{U}$. Then the following hold:*

(i) $\bigcap_{B \in SME_{\mathcal{W}}(A)} B = \mathrm{PSO}_{\mathcal{W}}(A)$*;*

(ii) $\mathrm{PSO}_{\mathcal{W}}(A)$ *is the set of all $\alpha \in A$ such that $A \setminus \{\alpha\} \not\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} \{\alpha\}$;*

(iii) $\mathrm{PSO}_{\mathcal{W}}(A) \equiv^{\mathcal{W}}_{\forall\forall\exists} A$ *if and only if $SME_{\mathcal{W}}(A)$ is a singleton, which is if and only if* $\mathrm{PSO}_{\mathcal{W}}(A)$ *is the unique setwise-minimal equivalent subset for A.*

4.2 Filtering

A simple way of generating a minimal equivalent subset of A is to sequentially delete elements α of A that are not needed for maintaining equivalence, i.e., are such that $A \setminus \{\alpha\} \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} \{\alpha\}$, since then $A \setminus \{\alpha\} \equiv^{\mathcal{W}}_{\forall\forall\exists} A$. This is what is done in the operation $\mathit{Filter}_\sigma(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists})$ defined below, to produce a minimal equivalent subset of A.[2]

For $\alpha \in A$, define $\mathit{Filter}(A, \alpha; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists})$ to be $A \setminus \{\alpha\}$ if $A \setminus \{\alpha\} \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} \{\alpha\}$; otherwise it equals A.

Let us label A as $\alpha_1, \dots, \alpha_n$, where $n = |A|$. Formally the labelling is a bijection σ from $\{1, \dots, n\}$ to A (so that $\sigma(i) = \alpha_i$), and let Λ be the set of all labellings. We define $\mathit{Filter}_\sigma(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists})$ iteratively as follows. We set $A^0 = A$. For $i = 1, \dots, n$, we set $A^i = \mathit{Filter}(A^{i-1}, \alpha_i; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists})$. We then define $\mathit{Filter}_\sigma(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists})$ to be A^n, i.e., the set remaining after iteratively deleting elements from A that are dominated w.r.t. relation $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$.

As the proposition below states, when applying the filtering operation $\mathit{Filter}_\sigma(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists})$, (i) equivalence is always maintained; and (ii) we always obtain a minimal equivalent subset, and any such subset can be achieved for some ordering. Part (iii) implies that for any labelling σ we have $\mathit{Filter}_\sigma(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}) = \mathrm{PSO}_{\mathcal{W}}(A)$ if $\mathrm{PSO}_{\mathcal{W}}(A) \equiv^{\mathcal{W}}_{\forall\forall\exists} A$.

Proposition 1 *Let $A \in \mathcal{M}$ and let σ be any labelling of A. Then we have:*

(i) $A \equiv^{\mathcal{W}}_{\forall\forall\exists} \mathit{Filter}_\sigma(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}) \subseteq A$.
(ii) $\mathit{SME}_{\mathcal{W}}(A) = \{\mathit{Filter}_\sigma(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}) : \sigma \in \Lambda\}$.
(iii) If $\mathrm{PSO}_{\mathcal{W}}(A) \equiv^{\mathcal{W}}_{\forall\forall\exists} A$ *then* $\mathit{Filter}_\sigma(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}) = \mathrm{PSO}_{\mathcal{W}}(A)$ *for any labelling* σ.

4.3 $\mathrm{PSO}_{\mathcal{W}}(A)$ as unique minimal equivalent set

We show that in certain very important classes of problem we do have $\mathrm{PSO}_{\mathcal{W}}(A) \equiv^{\mathcal{W}}_{\forall\forall\exists} A$, leading (by Theorem 1) to $\mathrm{PSO}_{\mathcal{W}}(A)$ being the unique setwise-minimal equivalent subset for equivalence-free A. The result below covers the linear case in which $f_w(\alpha) = w \cdot \alpha$, but also much more general forms of utility function. This contrasts with the general case in which $\mathrm{PSO}_{\mathcal{W}}(A)$ may well not be equivalent to A; it is even easy to construct small discrete examples in which $\mathrm{PSO}_{\mathcal{W}}(A)$ is empty; see e.g., Table 2 of [39].

Theorem 2 *Let $\Omega = \mathcal{U} = \mathbb{R}^p$ and let $\mathcal{W}$ be a convex subset of $\mathcal{U}$. Assume that for each $\alpha \in \Omega$, $\{f_w(\alpha) : w \in \mathcal{W}'\}$ is an analytic function of w, where $\mathcal{W}'$ is the smallest affine space containing $\mathcal{W}$. Assume that A ($\in \mathcal{M}$) is $\equiv_{\mathcal{W}}$-free. Then there exists a unique setwise-minimal equivalent subset for A, i.e., $\mathit{SME}_{\mathcal{W}}(A)$ is a singleton, and this equals $\mathrm{PSO}_{\mathcal{W}}(A)$, which equals the set of elements α of A such that $\mathrm{Opt}^A_{\mathcal{W}}(\alpha)$ has the same dimension as $\mathcal{W}$.*

5 SETWISE MAX REGRET

The condition $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ states that in every scenario, the set of alternatives A is at least as good as the set B. A natural related numerical measure is setwise max regret $\mathit{SMR}_{\mathcal{W}}(A, B)$, defined to be $\sup_{w \in \mathcal{W}} \mathit{Ut}_B(w) - \mathit{Ut}_A(w)$, which expresses how much worse A could be than B, i.e., the maximum regret of choosing A over B. When $A \subseteq B$, $\mathit{SMR}_{\mathcal{W}}(A, B)$ equals the setwise max regret $\mathit{SMR}(A, \mathcal{W})$ defined in [32]; that paper defines a method that involves finding a subset A of B (among a particular class of subsets, e.g., all those of a fixed cardinality k) that minimises $\mathit{SMR}_{\mathcal{W}}(A, B)$.[3] A can then be considered as a maximally informative query, to be used in an incremental elicitation process for finding an optimal element of B. $\mathit{SMR}_{\mathcal{W}}(A, B)$ is closely related also to the notion of setwise max regret defined in [4].

Regarding $\mathit{Ut}_A(w)$ as the utility achieved from set A in scenario w (and similarly, for $\mathit{Ut}_B(w)$), we have that $\mathit{SMR}_{\mathcal{W}}(A, B)$ is the worst-case loss of utility (or maximum regret) if we choose set A instead of set B. For instance if A is a subset of B, and $\mathit{SMR}_{\mathcal{W}}(A, B)$ is very close to zero, then we might consider that A is a sufficiently close approximation of B, simplifying the set of choices for the user. We have $\mathit{SMR}_{\mathcal{W}}(A, B) \leq 0$ if and only if $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ (see Proposition 2 below). The problem of computing $\mathit{SMR}_{\mathcal{W}}(A, B)$ is thus strongly related to that of determining $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$.

The definitions and results from earlier sections (apart from Section 4.3), regarding $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$, SME, PO, PSO and UD, depended only on the orderings $\succcurlyeq_w$, for $w \in \mathcal{W}$, and so were ordinal, in the sense that they are not affected by any strictly monotonic transformations of each function f_w (which can be different for each w). However, this is not the case for *SMR*, which has much weaker invariance properties.

We say that $\mathit{SMR}_{\mathcal{W}}(A, B)$ *is achieved* if there exists $w \in \mathcal{W}$ such that $\mathit{Ut}_B(w) - \mathit{Ut}_A(w) = \mathit{SMR}_{\mathcal{W}}(A, B)$, so that then $\mathit{SMR}_{\mathcal{W}}(A, B) = \max_{w \in \mathcal{W}} \mathit{Ut}_B(w) - \mathit{Ut}_A(w)$. We will mainly be interested in situations in which $\mathit{SMR}_{\mathcal{W}}(A, B)$ is achieved; this always happens, for instance, if for each $\alpha \in \Omega$, $f_w(\alpha)$ is a continuous function of w, and $\mathcal{W}$ is compact.

There are obvious monotonicity properties of $\mathit{SMR}_{\mathcal{W}}(A, B)$ with respect to $\mathcal{W}$, A and B. We give some further basic properties of the maximum regret function below. Parts (i) and (ii) give decomposability properties, with (i) being more useful computationally. (ii) is a slight generalisation of Observation 4 in [32]. (iii) relates the function $\mathit{SMR}_{\mathcal{W}}$ with the relation $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$, and (iv) with the Possibly Optimal operator $\mathrm{PO}_{\mathcal{W}}$, and (v) with the Possibly Strictly Optimal operator $\mathrm{PSO}_{\mathcal{W}}$. Property (vi) enables pre-processing of the sets A and B.

Proposition 2 *Consider $A, B \in \mathcal{M}$ and $\mathcal{W} \subseteq \mathcal{U}$.*

(i) $\mathit{SMR}_{\mathcal{W}}(A, B) = \max_{\beta \in B} \mathit{SMR}_{\mathcal{W}}(A, \{\beta\})$
(ii) $\mathit{SMR}_{\mathcal{W}}(A, B) = \max_{\alpha \in \mathrm{PO}(A)} \mathit{SMR}_{\mathrm{Opt}^A_{\mathcal{W}}(\alpha)}(\{\alpha\}, B)$.
(iii) $\mathit{SMR}_{\mathcal{W}}(A, B) \leq 0$ *if and only if* $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$.
(iv) If $\mathit{SMR}_{\mathcal{W}}(A, B)$ *is achieved then* $\mathit{SMR}_{\mathcal{W}}(A, B) \geq 0$ *if and only if* $\mathrm{PO}_{\mathcal{W}}(A \cup B) \cap B \neq \emptyset$.
(v) For equivalence-free A, and $\alpha \in A$, $\mathit{SMR}_{\mathcal{W}}(A \setminus \{\alpha\}, \{\alpha\}) > 0$ *if and only if* $\mathrm{PSO}_{\mathcal{W}}(A) \ni \alpha$.
(vi) If $A' \equiv^{\mathcal{W}}_{\forall\forall\exists} A$ *and* $B' \equiv^{\mathcal{W}}_{\forall\forall\exists} B$ *then* $\mathit{SMR}_{\mathcal{W}}(A', B') = \mathit{SMR}_{\mathcal{W}}(A, B)$.

6 IMPLICATIONS FOR INCREMENTAL PREFERENCE ELICITATION

In recent years there has been considerable focus in the AI preference community on incremental preference elicitation techniques, or

[2] One can define $\mathit{Filter}_\sigma(A; \gg^{\mathcal{W}}_{\forall\forall\exists})$ analogously, using the strongly strict version of the dominance relation, with the result being $\mathrm{PO}_{\mathcal{W}}(A)$, irrespective of σ, and without requiring any conditions on A; see Algorithm 1 and Theorem 1 in [2].

[3] In contrast, in generating $\mathrm{PSO}_{\mathcal{W}}(B)$ we are finding a minimal subset A of B such that $\mathit{SMR}_{\mathcal{W}}(A, B) = 0$ (under appropriate assumptions, such as those for Theorem 2).

active learning, see e.g., [14, 10, 32, 34, 2, 5, 7]. We argue that the notion of being possibly strictly optimal is important here.

Let α and β be alternatives. Preference model w is said to satisfy a preference statement $\alpha \geq \beta$ if $f_w(\alpha) \geq f_w(\beta)$, i.e., α is at least as good as β given w. For set of alternatives A the preference statement $\alpha \geq A$ means $\alpha \geq \beta$ for all $\beta \in A$. Thus, for $\alpha \in A$, w satisfies $\alpha \geq A$ if and only if (given w) α is a most preferred element in A, $\alpha \in \mathrm{O}_w(A)$, i.e., w makes α optimal in A. This holds if and only if $w \in \mathrm{Opt}^A_{\mathcal{W}}(\alpha)$.

In incremental elicitation a common strategy is to generate a small set of alternatives A, and to ask the user which element of A is most preferred. If they reply "α" then this is interpreted as $\alpha \geq A$. We will then update $\mathcal{W}$ to the set of all $w \in \mathcal{W}$ such that α is a most preferred option in A, i.e., we update $\mathcal{W}$ to $\mathrm{Opt}^A_{\mathcal{W}}(\alpha)$.

There can be forms of inconsistency, of different kinds, between the user answers and the model we have of the user. We say that, given set of preference models $\mathcal{W}$, alternative α is a *feasible answer to query* A if $\mathrm{Opt}^A_{\mathcal{W}}(\alpha)$ is non-empty, i.e., there exists some user preference model in $\mathcal{W}$ under which α is optimal in A.

We say that α is a *strongly feasible answer to query* A (given $\mathcal{W}$) if $\mathrm{Opt}^A_{\mathcal{W}}(\alpha)$ has the same dimension as $\mathcal{W}$. In the example in Figure 1, with the query $\{(11, 1), (10, 4), (7, 5), (6, 6), (4, 7)\}$, the elements $(11, 1)$ and $(7, 5)$ are infeasible answers, and $(6, 6)$ is not strongly feasible. The following result, which follows from Theorem 2, characterises [strongly] feasible answers to queries.

Proposition 3 *Consider $A \in \mathcal{M}$ and $\mathcal{W} \subseteq \mathcal{U}$.*

(i) α is a feasible answer to query A given $\mathcal{W}$ if and only if $\alpha \in \mathrm{PO}_{\mathcal{W}}(A)$.

(ii) Under the conditions of Theorem 2 on Ω, $\mathcal{U}$, $\mathcal{W}$ and f we have that α is a strongly feasible answer to query A given $\mathcal{W}$ if and only if $\alpha \in \mathrm{PSO}_{\mathcal{W}}(A)$.

If the user chooses α from A, and α is not a feasible answer to A, then we get an inconsistency, since the updated $\mathcal{W}$ will be empty. Suppose now, on the other hand, α is not a strongly feasible answer to A. We can still consistently update $\mathcal{W}$, so this is a less strong kind of inconsistency; however, such an answer would still be seriously troubling. For instance, suppose $\mathcal{W} \subseteq I\!R^p$, and consider any probability distribution over $\mathcal{W}$, regarding which is the true user model w, such that (as one would expect) the probability distribution is compatible with the measure of the sets. If α is not a strongly feasible answer to query A then the probability that w is such that $\alpha \geq A$ holds would be zero (since $\mathrm{Opt}^A_{\mathcal{W}}(\alpha)$ has then measure zero in $\mathcal{W}$, being of lower dimension than $\mathcal{W}$). A choice, by the user, of α would hence correspond with an event of probability zero.

To ensure that every answer to a query A is feasible, we thus require that $\mathrm{PO}_{\mathcal{W}}(A) = A$. And, to ensure that every answer to A is strongly feasible, we require that $\mathrm{PSO}_{\mathcal{W}}(A) = A$, i.e., that every element of A is strictly possibly optimal in A.

We thus argue that the standard methods for generating queries in incremental preference learning should be modified to ensure that every element in the query set is strictly possibly optimal.[4] Since Theorem 2 implies that $\mathrm{PSO}_{\mathcal{W}}(A)$ is non-empty, (and indeed equivalent to A) we can therefore replace a potential query A by $\mathrm{PSO}_{\mathcal{W}}(A)$.

It is shown in [32, 33] that choosing a subset A, of the set of available alternatives B, that maximises setwise regret $\mathit{SMR}_{\mathcal{W}}(A, B)$ (among small subsets) is a desirable and well-founded choice for an informative query. However, it can easily happen that, for such a query A, we have $\mathrm{PSO}_{\mathcal{W}}(A) \neq A$ and even $\mathrm{PO}_{\mathcal{W}}(A) \neq A$. Such a choice of A is then in danger of leading to an inconsistency, as described above. Fortunately, one can easily solve this problem by replacing A by $\mathrm{PSO}_{\mathcal{W}}(A)$, since if A maximises setwise regret then $\mathrm{PSO}_{\mathcal{W}}(A)$ also maximises setwise regret (under the conditions in Theorem 2 on Ω, $\mathcal{U}$, $\mathcal{W}$ and f) because $\mathit{SMR}_{\mathcal{W}}(\mathrm{PSO}_{\mathcal{W}}(A), B) = \mathit{SMR}_{\mathcal{W}}(A, B)$, by Theorem 2 and Proposition 2(vi).

[4] Learning an inconsistency could in theory be useful information, allowing the potential of updating the model in some way to restore consistency; however, this would probably have a heavy computational cost, and in a practical application, one will want to avoid the incremental elicitation procedure breaking down.

7 EEU METHOD FOR TESTING $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ AND COMPUTING $\mathit{SMR}_{\mathcal{W}}(A, B)$

Computing the extreme points of $\mathcal{W}$ can lead for the linear case to an easy way of testing if $\alpha \succcurlyeq_{\mathcal{W}} \beta$ (for $\alpha, \beta \in I\!R^p$): it is easy to see that $\alpha \succcurlyeq_{\mathcal{W}} \beta$ holds if and only if for each extreme point w of $\mathcal{W}$, we have $w \cdot (\alpha - \beta) \geq 0$ [22]. Similarly, it follows immediately that standard maximum regret over the convex polytope $\mathcal{W}$ can be computed using the extreme points of $\mathcal{W}$, as observed e.g., in [29]. However, for setwise max regret it is not sufficient to consider the extreme points of $\mathcal{W}$. Here we develop a novel extreme points method for testing $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ and computing $\mathit{SMR}_{\mathcal{W}}(A, B)$, by moving to a higher dimensional space.

Given $\mathcal{W}$, the utility function $\mathit{Ut}_A(w)$ (over $w \in \mathcal{W}$) can be viewed as a subset of $\mathcal{W} \times I\!R$, and we can test $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ by considering such subsets. Let us define $\Gamma(\mathcal{W}, A) \subseteq \mathcal{W} \times I\!R \subseteq I\!R^p \times I\!R$ to be $\{(w, r) : w \in \mathcal{W}, r \geq \mathit{Ut}_A(w)\}$, i.e., the *epigraph* [11] of the utility function Ut_A on $\mathcal{W}$. If $\mathcal{W}$ is convex and compact and for all $\alpha \in A$, $f_w(\alpha)$ is a convex and continuous function of $w \in \mathcal{W}$, then $\Gamma(\mathcal{W}, A)$ is a closed convex set. We write $\mathit{Ext}(\Gamma(\mathcal{W}, A))$ for the extreme points of $\Gamma(\mathcal{W}, A)$.

The following result leads to two different ways of testing $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$. Firstly, we can compute the extreme points of both $\Gamma(\mathcal{W}, A)$ and $\Gamma(\mathcal{W}, A \cup B)$; by (ii), these two sets of extreme points are equal if and only if $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$. Alternatively, we can test $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$, using part (iii), after computing $\mathit{Ext}(\Gamma(\mathcal{W}, A))$. We can compute the pairwise max regret $\mathit{SMR}_{\mathcal{W}}(A, B)$ as $\max_{\beta \in B} \mathit{SMR}_{\mathcal{W}}(A, \{\beta\})$ (see Proposition 2), and use part (iv) below.

Theorem 3 *Consider any finite subsets A and B of $I\!R^p$, any $\beta \in I\!R^p$, and any compact and convex subset $\mathcal{W}$ of $I\!R^p$, and assume that for all $\alpha \in A \cup B \cup \{\beta\}$, $f_w(\alpha)$ is a convex and continuous function of $w \in \mathcal{W}$.*

(i) $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B \iff \Gamma(\mathcal{W}, A) \subseteq \Gamma(\mathcal{W}, B) \iff \Gamma(\mathcal{W}, A) = \Gamma(\mathcal{W}, A \cup B)$.

(ii) $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ if and only if $\mathit{Ext}(\Gamma(\mathcal{W}, A)) = \mathit{Ext}(\Gamma(\mathcal{W}, A \cup B))$.

(iii) $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ holds if and only if for all $(w, r) \in \mathit{Ext}(\Gamma(\mathcal{W}, A))$ and for all $\beta \in B$ we have $f_w(\beta) \leq r$.

(iv) $\mathit{SMR}_{\mathcal{W}}(A, \{\beta\}) = \max\{f_w(\beta) - r : (w, r) \in \mathit{Ext}(\Gamma(\mathcal{W}, A))\}$.

Continuing the example, it can be seen from Figure 1 that $\mathit{Ext}(\Gamma(\mathcal{W}, A)) = \{(0, 7), (\frac{1}{3}, 6)(\frac{2}{3}, 8)\}$, where we are again abbreviating w to just its first component w_1, so that e.g., $(\frac{1}{3}, 6)$ represents the pair $(w, \mathit{Ut}_A(w))$ with $w = (\frac{1}{3}, \frac{2}{3})$. Then, using Theorem 3(iv), $\mathit{SMR}_{\mathcal{W}}(A, \{(5.5, 6.5)\}) = \max(-0.5, \frac{1}{6}, -2\frac{1}{6}) = \frac{1}{6} > 0$; for instance, the middle term in the max equals $f_w((5.5, 6.5)) - 6 = \frac{1}{3} \cdot 5.5 + \frac{2}{3} \cdot 6.5 - 6 = \frac{1}{6}$. Therefore, $A \not\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} \{(5.5, 6.5)\}$, by Proposition 2(iii). This illustrates the fact that it is not sufficient to just consider the extreme points of $\mathcal{W}$.

8 THE CASE OF MULTI-ATTRIBUTE UTILITY VECTORS

We now consider the situation in which we are especially interested, where the alternatives in Ω are multi-attribute utility vectors. Let $\mathcal{U}$ be the set of non-negative normalised vectors in $I\!R^p$, so that, for all $w \in \mathcal{U}$, for all $i = 1, \ldots, p$, $w_i \geq 0$, and $\sum_{i=1}^{p} w_i = 1$. (Thus, $\mathcal{U}$ is the *unit* $(p-1)$*-simplex*.) Let $\Omega = I\!R^p$. For $\alpha \in \Omega$ we define $f_w(\alpha) = \alpha \cdot w$, i.e., $\sum_{i=1}^{p} w_i \alpha_i$. This leads to, for $\alpha, \beta \in I\!R^p$, $\alpha \succcurlyeq_w \beta$ if and only if $(\alpha - \beta) \cdot w \geq 0$. Also, $Ut_A(w) = \max_{\alpha \in A} w \cdot \alpha$.

We will assume $\mathcal{W}$ to be a closed polytope in $I\!R^p$, which can be defined using a finite set of linear inequalities. Given a finite set $\Lambda = \{\lambda_i : i = 1, \ldots, k\}$ of vectors in $I\!R^p$, and corresponding real numbers r_i, we can define $\mathcal{W}$ to be the set of $w \in \mathcal{U}$ such that for all $i = 1, \ldots, k$, $w \cdot \lambda_i \geq r_i$. In particular, such linear inequalities can arise from input preferences of the form α *is preferred to* β, leading to the constraint $w \cdot (\alpha - \beta) \geq 0$.

This form of preferences has been studied a great deal; for instance, $\mathrm{UD}_{\mathcal{W}}(A)$ consists of the non-dominated alternatives in A for a multiobjective program (MOP) given a cone (with the cone generated as the dual of $\mathcal{W}$) [41, 37, 17]. Without any additional preferences, so that $\mathcal{W}$ is just the unit $(p-1)$-simplex, $\succcurlyeq_{\mathcal{W}}$ is the Pareto ordering on alternatives, and $\mathrm{UD}_{\mathcal{W}}(A)$ is set of Pareto-optimal alternatives, with the supported alternatives being also in $\mathrm{PO}_{\mathcal{W}}(A)$.

In Section 7 we gave an EEU method for computing $\mathit{SMR}_{\mathcal{W}}$ and testing dominance; in Section 8.1 we give a straight-forward LP method related to the approaches used in [32, 2, 4]. In Section 8.2 we give a result that enables one to compute the minimal equivalent subset using the extreme points of the epigraph.

8.1 Linear Programming for Computing $\mathit{SMR}_{\mathcal{W}}(A, B)$, and Testing $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$

$\mathit{SMR}_{\mathcal{W}}(A, \{\beta\})$ is equal to the maximum value of x such that there exists $w \in I\!R^p$ satisfying the constraints (i) $w \in \mathcal{W}$; and (ii) for all $\alpha \in A$, $w \cdot (\beta - \alpha) \geq x$. Since $\mathcal{W}$ is a closed polytope, we can use a linear programming solver to compute $\mathit{SMR}_{\mathcal{W}}(A, \{\beta\})$. Applying this for each $\beta \in A$ allows us to compute $\mathit{SMR}_{\mathcal{W}}(A, B)$, and thus, to test if $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$, using Proposition 2.

8.2 Using Extreme Points of Epigraph to Compute Minimal Equivalent Subset

For the linear case, with $\alpha \in A$, the set $\mathrm{Opt}^A_{\mathcal{W}}(\alpha)$, consisting of all w in $\mathcal{W}$ that make α optimal in A (see Section 4.1), is convex; we abbreviate its set of extreme points $\mathit{Ext}(\mathrm{Opt}^A_{\mathcal{W}}(\alpha))$ to $E^A_{\mathcal{W}}(\alpha)$. Theorem 2 implies that $\mathrm{PSO}_{\mathcal{W}}(A)$ is the unique minimal equivalent subset of an (equivalence-free) set $A \in \mathcal{M}$, which can be shown to consist of all $\alpha \in A$ such that there does not exist $\beta \in A$ such that $\mathrm{Opt}^A_{\mathcal{W}}(\beta) \supsetneqq \mathrm{Opt}^A_{\mathcal{W}}(\alpha)$. The following result shows that the condition $\mathrm{Opt}^A_{\mathcal{W}}(\beta) \supsetneqq \mathrm{Opt}^A_{\mathcal{W}}(\alpha)$ is (perhaps surprisingly) equivalent to $E^A_{\mathcal{W}}(\beta) \supsetneqq E^A_{\mathcal{W}}(\alpha)$, and that $E^A_{\mathcal{W}}(\alpha)$ can be computed by projecting the extreme points of the epigraph, $\mathit{Ext}(\Gamma(\mathcal{W}, A))$. This is the basis of our method, described in Section 9.1(II) below, for efficiently computing the minimal equivalent set $\mathrm{PSO}_{\mathcal{W}}(A)$.

Proposition 4 *Assume that $\mathcal{W}$ is a convex subset of $I\!R^p$, and that for $w \in I\!R^p, \alpha \in I\!R^p$, $f_w(\alpha) = w \cdot \alpha$. Consider $A \in \mathcal{M}$, $w \in \mathcal{W}$, and $\alpha, \beta \in A$.*

(i) $\mathrm{Opt}^A_{\mathcal{W}}(\alpha) \subseteq \mathrm{Opt}^A_{\mathcal{W}}(\beta) \iff E^A_{\mathcal{W}}(\alpha) \subseteq E^A_{\mathcal{W}}(\beta)$.
(ii) $E^A_{\mathcal{W}}(\alpha) = \{w \in I\!R^p : (w, w \cdot \alpha) \in \mathit{Ext}(\Gamma(\mathcal{W}, A))\}$.
(iii) If $\mathcal{W}$ is compact then $\mathit{dim}(\mathrm{Opt}^A_{\mathcal{W}}(\alpha)) < |E^A_{\mathcal{W}}(\alpha)|$.

9 THE STRUCTURE OF THE ALGORITHMS

In this section we make use of mathematical results in previous sections in developing computational methods for computing the minimal equivalent set $\mathrm{PSO}_{\mathcal{W}}(A)$ and testing dominance between sets, for the case of multi-attribute utility vectors, with the set of scenarios $\mathcal{W}$ being a convex polytope, and with linear utility functions.

9.1 Computing Minimal Equivalent Set

Given $A \in \mathcal{M}$, we aim to generate $A' \subseteq A$ with $A' \equiv^{\mathcal{W}}_{\forall\forall\exists} A$, and such that for strict subset A'' of A', $A'' \not\equiv^{\mathcal{W}}_{\forall\forall\exists} A$.

First we pre-process by eliminating elements of A not in $\mathrm{UD}_{\mathcal{W}}(A)$. At the same time we can make A equivalence-free. Theorem 2 implies that there exists a minimal equivalent set, i.e., $\mathrm{SME}_{\mathcal{W}}(A)$ has a unique element, say, A', and this equals $\mathrm{PSO}_{\mathcal{W}}(A)$. We have two methods for then computing A'.

(I) we use the approach $\mathit{Filter}_{\sigma}(A; \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists})$ defined in Section 4.2. This involves multiple (i.e., $|A|$) tests of the form $A \setminus \{\alpha\} \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} \{\alpha\}$, which can be achieved using a similar approach to 9.2 below, using an LP solver.

(II) For each $\alpha \in A$ we compute $E^A_{\mathcal{W}}(\alpha)$ using Proposition 4(ii), by computing the extreme points of the epigraph. We can eliminate any element α such that $|E^A_{\mathcal{W}}(\alpha)| \leq \mathit{dim}(\mathcal{W})$, since Proposition 4(iii) would then imply that $\mathit{dim}(\mathrm{Opt}^A_{\mathcal{W}}(\alpha)) < \mathit{dim}(\mathcal{W})$, and thus, $\alpha \notin \mathrm{PSO}_{\mathcal{W}}(A)$, by Theorem 2. If for any $\alpha, \beta \in A$ with $\alpha \not\equiv_{\mathcal{W}} \beta$, $E^A_{\mathcal{W}}(\alpha) \subseteq E^A_{\mathcal{W}}(\beta)$, (so then $\mathrm{Opt}^A_{\mathcal{W}}(\alpha) \subseteq \mathrm{Opt}^A_{\mathcal{W}}(\beta)$) then we know that $\alpha \notin \mathrm{PSO}_{\mathcal{W}}(A)$, and if those sets are equal we know also that $\beta \notin \mathrm{PSO}_{\mathcal{W}}(A)$. Now, any of these elements can then be deleted from A, because of an incrementality property of $\mathrm{PSO}_{\mathcal{W}}$. We then continue until for all remaining elements $\alpha, \beta \in A$ we have $E^A_{\mathcal{W}}(\alpha) \not\subseteq E^A_{\mathcal{W}}(\beta)$, and then $A = A'$, the unique element of $\mathrm{SME}_{\mathcal{W}}(A)$, the set of possibly strictly optimal elements.

9.2 Testing $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$

Our algorithm includes three steps of increasing complexity:

(1) Efficiently testing (a) a necessary condition $A \succcurlyeq^{\mathcal{W}^0}_{\forall\forall\exists} B$, where $\mathcal{W}^0 = \mathit{Ext}(\mathcal{W})$ is the set of extreme points of $\mathcal{W}$; and (b) a sufficient condition, whether there exists $\alpha \in A$ such that for all $w \in \mathcal{W}^0$, $f_w(\alpha) \geq Ut_B(w)$; (the conditions can be tested together, by first computing $Ut_B(w)$ for each $w \in \mathcal{W}^0$). If (a) is false then we know that $A \not\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ (because of monotonicity with respect to $\mathcal{W}$); if (b) is true then we know that $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ holds. If the necessary condition is false, or the sufficient condition is true, then we need go no further. The complexity of this step is proportional to $|A| + |B|$.

(2) Pre-processing by reducing the sets A and B; this step has complexity proportional to $|A||B|$. We replace A by $\mathrm{UD}_{\mathcal{W}}(A)$ and B by $\mathrm{UD}_{\mathcal{W}}(B)$. We then eliminate all elements β from B such that for some $\alpha \in A$, $\alpha \succcurlyeq_{\mathcal{W}} \beta$. If B becomes empty then we can stop, since we then have $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$.

(3) We determine whether $A \succcurlyeq^{\mathcal{W}}_{\forall\forall\exists} B$ holds using one of the methods in Sections 8.1 and 7, i.e., doing either (a), (b) or (c) below:

   (a) Using linear programming, as described in Section 8.1.

   (b) Using Theorem 3(ii) and testing if $\mathit{Ext}(\Gamma(\mathcal{W}, A))$ equals $\mathit{Ext}(\Gamma(\mathcal{W}, A \cup B))$.

(c) Using Theorem 3(iii) and testing if $f_w(\beta) \leq r$ for all $\beta \in B$ and for all $(w, r) \in Ext(\Gamma(\mathcal{W}, A))$.

Although we focus on non-strict dominance, the same algorithms can also be used to test the strongly strict dominance $A \gg^{\mathcal{W}}_{\forall\exists} B$ given as for all $w \in \mathcal{W}$, $Ut_A(w) > Ut_B(w)$. In particular, under the conditions of Theorem 3, we have $A \gg^{\mathcal{W}}_{\forall\exists} B \iff SMR_{\mathcal{W}}(A, B) < 0$, which is if and only if for all $(w, r) \in Ext(\Gamma(\mathcal{W}, A))$ and for all $\beta \in B$ we have $f_w(\beta) < r$.

10 EXPERIMENTAL TESTING

We briefly summarise the results of our experimental testing; more details are included in the longer version [30]. All experiments were performed on a computer facilitated by a Core i5 2.70 GHz processor and 8 GB RAM. We used CPLEX 12.8 [21] as the linear programming solver, and we used the Python library pycddlib [31] for computing the extreme points of a polytope.

We consider the linear case, where $\mathcal{W}$ is a subset of the unit $(p-1)$-simplex which is an intersection of T half-spaces. Specifically, we choose T (consistent) random user preferences of the form $aw_i + bw_j \geq cw_k$ (meaning that the user prefers a units of w_i and b units of w_j to c units of w_k), like in [24]. The alternatives in the sets A and B are integer utility vectors.

The pre-processing steps based on the $\mathrm{UD}_{\mathcal{W}}$ filtering were very worthwhile, for both computing the minimal set in Section 9.1, and in 9.2(2) for dominance; they reduce the sizes of the sets A and B very considerably (see e.g., Table 2), making the algorithms much faster overall, e.g., by an order of magnitude.

For cases in which $dim(\mathcal{W}) < 7$, the EEU method 9.1(I) to compute $\mathrm{PSO}_{\mathcal{W}}(A)$ was on average faster, and scaled better with the cardinalities of sets A and B, than the LP method 9.1(II). However, the situation dramatically reverses for $dim(\mathcal{W}) \geq 7$; this may well be because the number of extreme points is much larger then. This is illustrated by Table 1 along with the performance of the $\mathrm{UD}_{\mathcal{W}}$ filtering, where each figure is an average over 100 random instances. We also tested our algorithms with larger A such as $|A| =$ 20,000, with $dim(\mathcal{W}) = 5$ and four user preferences giving an average execution time over 100 experiments of 13 seconds for the $\mathrm{UD}_{\mathcal{W}}$ filtering, 22 seconds for the LP-based method and 6 seconds for the EEU method.

$dim(\mathcal{W})$	$\mathrm{UD}_{\mathcal{W}}$[s]	LP[s]	EEU[s]	# extreme points
2	0.014	0.057	0.001	13.24
3	0.036	0.192	0.005	57.52
4	0.116	0.548	0.039	248.14
5	0.268	1.467	0.310	1024.74
6	0.439	3.062	2.103	3667.36
7	0.683	5.943	15.630	13483.87

Table 1. Execution times (in seconds) of methods to compute $\mathrm{PSO}_{\mathcal{W}}(A)$ (Section 9.1), i.e., $\mathrm{UD}_{\mathcal{W}}$ filtering, EEU (I) and LP (II) (and number of extreme points of the epigraph), w.r.t. $dim(\mathcal{W})$ with $|A| = 500$ and 4 user preferences.

We also tested our EEU approach against the standard LP approach to compute $\mathrm{PO}_{\mathcal{W}}$ and also in this case it looks like that EEU is faster for cases in which $dim(\mathcal{W}) \leq 6$. The performances of EEU to compute $\mathrm{PO}_{\mathcal{W}}$ are very similar to the performances of EEU to compute $\mathrm{PSO}_{\mathcal{W}}$ shown in Table 1. With sets generated with our random problem generator we observed that the $\mathrm{PSO}_{\mathcal{W}}$ filtering removes around 5% more elements than $\mathrm{PO}_{\mathcal{W}}$.

Tables 2 and 3 give results for testing $A \succcurlyeq^{\mathcal{W}}_{\forall\exists} B$ (Section 9.2), where each figure is an average over 100 instances in which the initial test 9.2(1) was inconclusive (i.e., failed to determine whether or not $A \succcurlyeq^{\mathcal{W}}_{\forall\exists} B$ holds), and the size of the set B after the $\mathrm{UD}_{\mathcal{W}}$ filtering 9.2(2) was greater than zero.

Table 2 shows how the input sets A and B were reduced by the $\mathrm{UD}_{\mathcal{W}}$ filtering 9.2 (2). As we can see, increasing the size of $dim(\mathcal{W})$, the number of undominated elements increase and therefore the number of elements removed by the $\mathrm{UD}_{\mathcal{W}}$ filtering reduces.

$dim(\mathcal{W})$	$\|A'\|$	$\|B'\|$	$\|B''\|$
2	10.08	7.99	2.44
3	23.89	22.04	5.11
4	48.70	19.53	11.93
5	92.23	86.94	15.45
6	144.78	142.89	41.15
7	206.93	210.11	64.28

Table 2. Number of elements of $A' = \mathrm{UD}_{\mathcal{W}}(A)$, $B' = \mathrm{UD}_{\mathcal{W}}(B)$ and $B'' = \{\beta \in B' : \forall \alpha \in A,\ \alpha \not\succcurlyeq_{\mathcal{W}} \beta\}$ w.r.t. $dim(\mathcal{W})$ with $|A| = |B| = 500$ and 4 user preferences.

Table 3 gives average execution time of the preliminaries steps and the methods 3(a), 3(b) and 3(c) of Section 9.2. The checking of the necessary and the sufficient condition in 9.2(1) were very effective: on approximately 94% of the problems generated with our random problem generator, the necessary condition failed, or the sufficient condition succeeded, allowing the algorithm to stop in advance. On average, method 3(c) seems to be faster than method 3(b), and the LP method seems to the fastest for $dim(\mathcal{W}) \geq 6$. As for the previous case, the EEU methods are much worse for the case of $dim(\mathcal{W}) = 7$.

$dim(\mathcal{W})$	NSc[s]	$\mathrm{UD}_{\mathcal{W}}$[s]	$\mathrm{T_{LP}}$[s]	$\mathrm{T_{EPU}}$[s]	$\mathrm{T_{EEU}}$[s]
2	0.008	0.028	0.016	0.001	0.001
3	0.010	0.070	0.038	0.003	0.002
4	0.013	0.223	0.131	0.015	0.013
5	0.015	0.504	0.238	0.105	0.088
6	0.015	0.858	0.967	1.179	1.028
7	0.016	1.487	2.232	24.452	14.97

Table 3. Execution time of methods for testing the dominance $A \succcurlyeq^{\mathcal{W}}_{\forall\exists} B$ (Section 9.2), i.e., testing the necessary and the sufficient condition (NSc) (1), $\mathrm{UD}_{\mathcal{W}}$ filtering (2), $\mathrm{T_{LP}}$ 3(a) and $\mathrm{T_{EPU}}$ 3(b), $\mathrm{T_{EEU}}$ 3(c) for testing $A \succcurlyeq^{\mathcal{W}}_{\forall\exists} B$, w.r.t. $dim(\mathcal{W})$ with $|A| = |B| = 500$ and 4 user preferences.

11 DISCUSSION

We defined natural notions of equivalence and dominance for a general model of sets of multi-attribute utility, and proved general properties. Computationally we focused especially on the linear (weighted sum) case and we proved that there is a unique setwise-minimal equivalent subset of any (equivalence-free) set of utility vectors A. This set then equals the set of possibly strictly optimal alternatives $\mathrm{PSO}(A)$, and is a compact representation of the utility function for A, giving the utility achievable with A for each scenario. We show that filtering a query with the PSO operator avoids the potential of inconsistency in the user response. Along with preprocessing techniques we developed a linear programming method for generating $\mathrm{PSO}(A)$, and a method based on computing the extreme points of the epigraph of the utility function (EEU), as well as related methods for testing dominance. We implemented the approaches and our testing on random problems showed that both methods scaled to substantially sized problems, with the EEU method being better for lower dimensions. Our methods can be directly applied to reduce the set of utility vectors derived for a multi-objective influence diagram [24] or a multi-objective optimisation problem [25]. In the latter cases, A is derived from a combinatorial structure; in the future we plan to consider further developments of the computational techniques that make use of such combinatorial structures.

Although we focus on the case where $f_w(\alpha)$ is linear in w, it would be natural to develop computational procedures for non-linear cases (perhaps e.g., [6]), based on our more general characterisation results, such as Theorems 2 and 3 and Proposition 2.

A further natural application of our model and methods is for computing the Value of Information [16] for a multi-objective influence diagram. Each observable variable generates a Value of Information function which is a utility function Ut_A, so different observable variables can be compared using the relation $\succcurlyeq^{\mathcal{W}}_{\forall\forall\exists}$.

ACKNOWLEDGEMENTS

This material is based upon works supported by the Science Foundation Ireland under Grant No. 12/RC/2289 and Grant No. 12/RC/2289-P2 which are co-funded under the European Regional Development Fund.

REFERENCES

[1] Antreas D Athanassopoulos and Victor V Podinovski, 'Dominance and potential optimality in multiple criteria decision analysis with imprecise information', *Journal of the Operational research Society*, **48**(2), 142–150, (1997).
[2] N. Benabbou and P. Perny, 'Incremental weight elicitation for multiobjective state space search', in *Proc. AAAI 2015*, pp. 1093–1099, (2015).
[3] N. Benabbou and P. Perny, 'On possibly optimal tradeoffs in multicriteria spanning tree problems', in *Proc. ADT 2015*, volume 9346 of *Lecture Notes in Computer Science*, pp. 322–337. Springer, (2015).
[4] N. Benabbou and P. Perny, 'Solving multi-agent knapsack problems using incremental approval voting', in *Proc. ECAI-2016*, pp. 1318–1326. IOS Press, (2016).
[5] N. Benabbou and P. Perny, 'Adaptive elicitation of preferences under uncertainty in sequential decision making problems', in *Proc. IJCAI 2017*, pp. 4566–4572, (2017).
[6] N. Benabbou, P. Perny, and P. Viappiani, 'Incremental elicitation of Choquet capacities for multicriteria choice, ranking and sorting problems', *Artif. Intell.*, **246**, 152–180, (2017).
[7] N. Bourdache and P. Perny, 'Active preference elicitation based on generalized gini functions: Application to the multiagent knapsack problem', in *Proc. AAAI-2019*, (2019).
[8] Nadjet Bourdache, Patrice Perny, and Olivier Spanjaard, 'Incremental elicitation of rank-dependent aggregation functions based on bayesian linear regression', in *Proc. IJCAI2019*, pp. 2023–2029, (2019).
[9] Craig Boutilier, 'A POMDP formulation of preference elicitation problems', in *AAAI/IAAI*, pp. 239–246, (2002).
[10] Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans, 'Constraint-based optimization and utility elicitation using the minimax decision criterion', *Artificial Intelligence*, **170**(8-9), 686–713, (2006).
[11] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, England, 2004.
[12] Darius Braziunas and Craig Boutilier, 'Minimax regret based elicitation of generalized additive utilities.', in *UAI*, pp. 25–32, (2007).
[13] D. G. Bridge and F. Ricci, 'Supporting product selection with query editing recommendations', in *Proc. RecSys-2007*, pp. 65–72. ACM, (2007).
[14] U. Chajewska, D. Koller, and R. Parr, 'Making rational decisions using adaptive utility elicitation', in *Proc. AAAI2000*, pp. 363–369, (2000).
[15] M. Diehl and Y. Haimes, 'Influence diagrams with multiple objectives and tradeoff analysis', *IEEE Transactions On Systems, Man, and Cybernetics Part A*, **34**(3), 293–304, (2004).
[16] S. L. Dittmer and F. V. Jensen, 'Myopic value of information in influence diagrams', in *Proc. UAI '97:*, pp. 142–149, (1997).
[17] M. Ehrgott and M. M. Wiecek, *Mutiobjective Programming*, 667–708, Springer New York, New York, NY, 2005.
[18] M. Gelain, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh, 'Elicitation strategies for soft constraint problems with missing preferences: Properties, algorithms and experimental studies', *Artif. Intell.*, **174**(3-4), 270–294, (2010).
[19] Mirco Gelain, Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, and Nic Wilson, 'Interval-valued soft constraint problems', *Ann. Math. Artif. Intell.*, **58**(3-4), 261–298, (2010).
[20] Gordon B Hazen, 'Partial information, dominance, and potential optimality in multiattribute utility theory', *Operations research*, **34**(2), 296–310, (1986).
[21] IBM ILOG. IBM ILOG CPLEX Optimization Studio, V12.8.0, 2017.
[22] S. Kaddani, D. Vanderpooten, J. M. Vanpeperstraete, and H. Aissi, 'Weighted sum model with partial preference information: Application to multi-objective optimization', *European Journal of Operational Research*, **260**(2), 665–679, (2017).
[23] R. Keeney and H. Raiffa, *Decisions with multiple objectives: preferences and value tradeoffs*, Cambridge University Press, 1993.
[24] R. Marinescu, A. Razak, and N. Wilson, 'Multi-objective influence diagrams', in *Proc. UAI-2012*, pp. 574–583, (2012).
[25] R. Marinescu, A. Razak, and N. Wilson, 'Multi-objective constraint optimization with tradeoffs', in *Proc. CP-2013*, pp. 497–512, (2013).
[26] R. Marinescu, A. Razak, and N. Wilson, 'Multi-objective influence diagrams with possibly optimal policies', in *Proc. AAAI-2017*, pp. 3783–3789, (2017).
[27] C. O'Mahony and N. Wilson, 'Sorted-pareto dominance and qualitative notions of optimality', in *Proc. ECSQARU'2013*, pp. 449–460, (2013).
[28] A. Salo and R. P. Hämäläinen, 'Preference programming – multicriteria weighting models under incomplete information', in *Handbook of Multicriteria Analysis*, pp. 167–187. Springer Berlin Heidelberg, (2010).
[29] Mikhail Timonin, 'Robust optimization of the Choquet integral', *Fuzzy Sets and Systems*, **213**, 27–46, (2013).
[30] Federico Toffano and Nic Wilson, *Minimality and Comparison of Sets of Multi-Attribute Vectors (Longer Version)*, http://ucc.insight-centre.org/nwilson/MinimalitySetsLonger.pdf, 2019.
[31] Matthias C. M. Troffaes. pycddlib Python wrapper for Komei Fukuda's cddlib, https://pycddlib.readthedocs.io/en/latest/, 2018.
[32] P. Viappiani and C. Boutilier, 'Regret-based optimal recommendation sets in conversational recommender systems', in *Proc. RecSys-2009*, pp. 101–108. ACM, (2009).
[33] P. Viappiani and C. Boutilier, 'Recommendation sets and choice queries: There is no exploration/exploitation tradeoff!', in *Proc. AAAI-2011*. AAAI Press, (2011).
[34] Paolo Viappiani and Craig Boutilier, 'Optimal Bayesian recommendation sets and myopically optimal choice query sets', in *Proc. NIPS 2010*, pp. 2352–2360, (2010).
[35] Martin Weber, 'Decision making with incomplete information', *European journal of operational research*, **28**(1), 44–57, (1987).
[36] C. White, A. P. Sage, and S. Dozono, 'A model of multiattribute decision-making and trade-off weight determination under uncertainty', *IEEE Transactions on Systems, Man, and Cybernetics*, **14**(2), 223–229, (1984).
[37] M. Wiecek, 'Advances in cone-based preference modeling for decision making with multiple criteria', *Decision Making in Manufacturing and Services*, **1**(1-2), 153–173, (2007).
[38] N. Wilson and A.-M. George, 'Efficient inference and computation of optimal alternatives for preference languages based on lexicographic models', in *Proc. IJCAI 2017*, pp. 1311–1317, (2017).
[39] N. Wilson and C. O'Mahony, 'The relationships between qualitative notions of optimality for decision making under logical uncertainty', in *Proc. AICS-2011*, (2011).
[40] N. Wilson, A. Razak, and R. Marinescu, 'Computing possibly optimal solutions for multi-objective constraint optimisation with tradeoffs', in *Proc. IJCAI-2015*, (2015).
[41] P. Yu, 'Cone convexity, cone extreme points, and nondominated solutions in decision problems with multiobjectives', *Journal of Optimization Theory and Applications*, **14**(3), 319–377, (1974).