Edges of Mutually Non-dominating Sets

Richard M. Everson
Computer Science
University of Exeter, UK
R.M.Everson@ex.ac.uk

David J. Walker
Computer Science
University of Exeter, UK
D.J.Walker@ex.ac.uk

Jonathan E. Fieldsend
Computer Science
University of Exeter, UK
J.E.Fieldsend@ex.ac.uk

ABSTRACT

Multi-objective optimisation yields an estimated Pareto front of mutually non-dominating solutions, but with more than three objectives understanding the relationships between solutions is challenging. Natural solutions to use as *landmarks* are those lying near to the edges of the mutually non-dominating set. We propose four definitions of *edge points* for many-objective mutually non-dominating sets and examine the relations between them.

The first defines edge points to be those that extend the range of the attainment surface. This is shown to be equivalent to finding points which are not dominated on projection onto subsets of the objectives. If the objectives are to be minimised, a further definition considers points which are not dominated under maximisation when projected onto objective subsets. A final definition looks for edges via alternative projections of the set.

We examine the relations between these definitions and their efficacy for synthetic concave- and convex-shaped sets, and on solutions to a prototypical many-objective optimisation problem, showing how they can reveal information about the structure of the estimated Pareto front.

Categories and Subject Descriptors

G.1.6 [**Mathematics of computing**]: Global optimization; G.1.10 [**Mathematics of computing**]: Applications

Keywords

Mutually non-dominating sets; many-objective optimisation; visualisation; edge; extrema

1. INTRODUCTION

As optimisation algorithms become capable of tackling multi-objective problems with at least four objectives, often called many-objective problems, it becomes important to find ways of understanding and visualising the solutions in the approximation to the many-objective Pareto front [1].

GECCO'13, July 6–10, 2013, Amsterdam, The Netherlands.

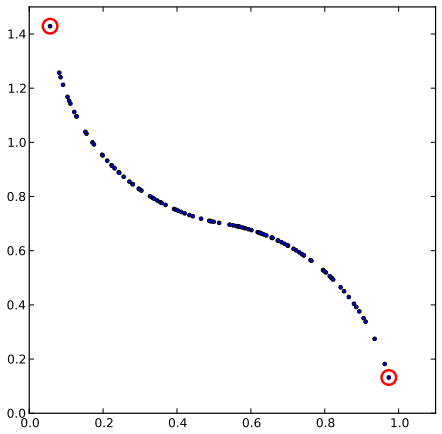

Figure 1: A two objective mutually non-dominating set with the two extremal individuals, corresponding to the edges, marked with circles.

Recently a number of methods have been developed to map the mutually non-dominated set of solutions to the plane or three dimensions for visualisation [2, 3]. However, the inevitable loss of information concomitant with the dimension reduction means that the problem owner or decision maker may need help in navigating the visualisation. One way of doing that is to identify *landmark* solutions — individuals with known properties — against which other individuals may be compared. Natural landmark individuals are extremal in one sense or another. For example, the individuals that maximise or minimise any single objective provide natural reference points [3]. Singh *et al.* [4] have given a procedure for finding the *corners* in multi-objective optimisation problems. In this paper we extend these ideas by examining what is meant by the *edge* of a mutually non-dominating set.

Although we concentrate on the visualisation aspects of edge points, they are also of interest for the design of multi-objective evolutionary algorithms, where it is useful to preferentially retain edge points in a search population, because these points preserve the spread of the search and are inherently diverse (e.g., [5, 6, 7]).

With two objectives the idea of edges is intuitively straightforward: as illustrated in Figure 1, the two individuals lying at the ends of the set comprise the edges and all the other individuals lie in the interior. We assume throughout that the criteria are to be minimised. Likewise, with three objectives

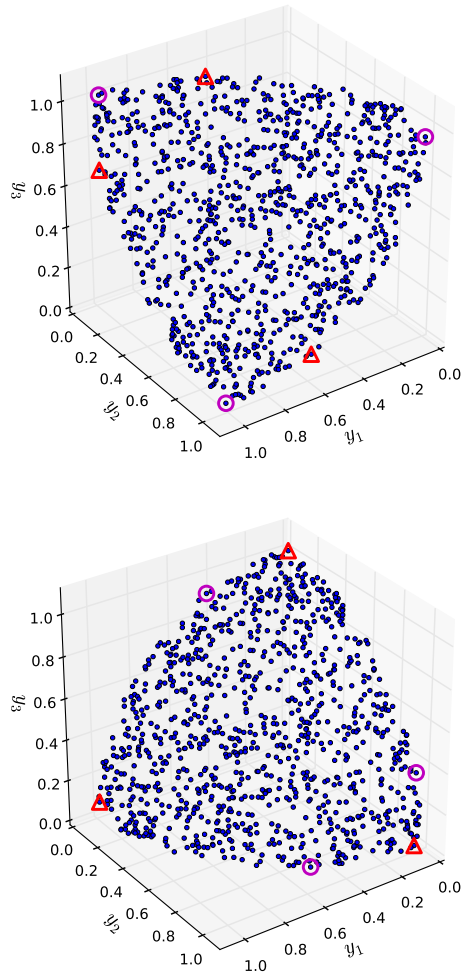

Figure 2: Three objective concave and convex mutually non-dominating sets showing extremal individuals. Individuals minimising and maximising objectives are marked with magenta circles and red triangles respectively. The convex set comprises 1000 points uniformly sampled from the positive octant a shell of unit radius. The concave set is generated from the positive set by $\mathbf{y} \mapsto -\mathbf{y} + (1, 1, 1)^T$.

the observer of a scatter plot of a mutually non-dominating set such as that shown in Figure 2 can readily identify which points are close to the edges of the set and which points are in the interior. Nonetheless, as we discuss later, defining precisely which individuals comprise the edges in this case is not entirely straightforward. Identifying these edges is important for understanding the extent of the set, although it is generally unlikely that a decision maker, in choosing one particular solution from the set, will pick one of the edge individuals, rather preferring another solution that makes a trade-off between all the objectives. With more than three objectives the edge points cannot be directly identified visually and it is not *a priori* clear that visualisation methods that map the set to the plane also map the edges in the original high-dimensional objective space to the edges of the planar visualisation.

In the rest of this paper we examine four definitions of the

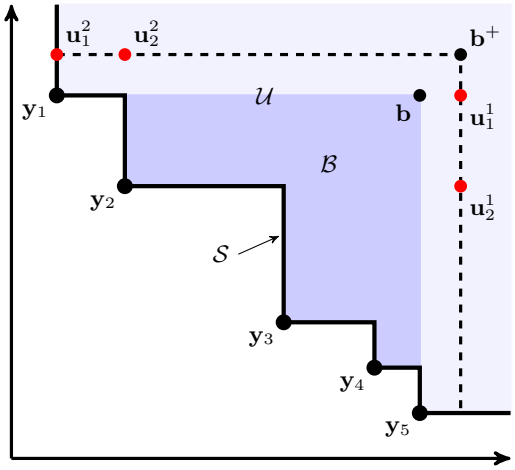

Figure 3: A mutually non-dominating set of points $\mathcal{Y} = \{\mathbf{y}_i\}$ and the region that it dominates $\mathcal{U}$. The attainment surface $\mathcal{S}$ is shown with the thick black line. Points in the region $\mathcal{B}$ dominate the reference point $\mathbf{b}$, whose mth coordinate is $b_m = \max_{\mathbf{y}' \in \mathcal{Y}} y'_m$. Candidate points for sole domination by $\mathbf{y}_1$ ($\mathbf{u}_1^1$ and $\mathbf{u}_1^2$) and $\mathbf{y}_2$ ($\mathbf{u}_2^1$ and $\mathbf{u}_2^2$) are shown in red; see equation (6).

edges of a mutually non-dominating set, all of which generalise the straightforward two-objective notion. Briefly, the first of these defines edge points as those that extend the range of the attainment surface. This turns out to identify the same edge individuals as a straightforward generalisation of Singh *et al.*'s corner points. The other two possible definitions look for edges in low-dimensional projections of the set.

We are particularly concerned with the performance of these definitions for many-objective sets. In addition to synthetic two and three dimensional datasets, which are useful because the edge points are easily visualised, we therefore examine these methods on set of solutions to a 9-objective problem, proposed as a prototypical many-objective problem by Hughes [8, 9]. Hughes considered the problem of designing an appropriate set of waveforms that can be transmitted by a Pulsed Doppler Radar to simultaneously measure the velocity and distance of a target. In the data we use all the objectives are to be minimised. Objectives 1, 3, 5, and 7 are concerned with range, while objectives 2, 4, 6 and 8 relate to velocity; objective 9 is the total transmission time of the waveform which is being optimised.

To be definite we assume that the set of mutually non-dominating individuals $\mathcal{Y} = \{\mathbf{y}_n\}_{n=1}^N$ comprises N individuals $\mathbf{y}_n$, each of which is an M-dimensional vector of objective values; $y_{n,m}$ denotes the value of the mth objective for the nth point. Without loss of generality, we assume that minimising the objective values corresponds to good performance.

2. ATTAINMENT SURFACE EDGES

The idea underlying this characterisation of the edge points of $\mathcal{Y}$ is that they extend the attainment surface [10]; that is, they dominate (unbounded) regions that are not dominated by other elements of $\mathcal{Y}$. This is illustrated for a two-dimensional set in Figure 3, in which the region $\{(u_1, u_2) \,|\, y_{11} <$

$u_1 \leq y_{21}$ and $y_{12} < u_2\}$ is dominated only by $\mathbf{y}_1$. This region is unbounded, but note that there are regions close to the attainment surface which are also dominated by a single element of $\mathcal{Y}$, but are bounded. For example, only $\mathbf{y}_3$ dominates the rectangle $\{(u_1, u_2) \,|\, y_{31} < u_1 < y_{41}$ and $y_{32} < u_2 < y_{22}\}$. Our definition of an edge point of $\mathcal{Y}$ is thus a point which extends the region dominated by the attainment surface by appending an unbounded region.

To make this precise define the region dominated by $\mathcal{Y}$:

$$\mathcal{U} = \{\mathbf{y} \,|\, \mathbf{u} \prec \mathbf{y} \text{ for some } \mathbf{u} \in \mathcal{Y}\} \tag{1}$$

The attainment surface is the boundary of $\mathcal{U}$ [7]. Also, let $\mathbf{b}$ be the point which has the maximum coordinate of any element of $\mathcal{Y}$ in each dimension:

$$b_m = \max_{\mathbf{y}_n \in \mathcal{Y}} y_{nm} \tag{2}$$

Then let $\mathcal{B}$ be the region which is dominated by $\mathcal{Y}$ but lies "below" $\mathbf{b}$:

$$\mathcal{B} = \mathcal{U} \cap \{\mathbf{y} \,|\, \mathbf{y} \prec \mathbf{b}\} \tag{3}$$

Finally, we define a function that returns all the elements of $\mathcal{Y}$ that (weakly) dominate a point $\mathbf{y}$:

$$\text{doms}_{\mathcal{Y}}(\mathbf{y}) = \{\mathbf{u} \in \mathcal{Y} \,|\, \mathbf{u} \preceq \mathbf{y}\} \tag{4}$$

Given these preliminaries, we define $\mathbf{y}$ to be an *attainment surface edge point* of $\mathcal{Y}$ if and only if there exists a $\mathbf{u} \in \mathcal{U} \setminus \mathcal{B}$ such that $|\,\text{doms}_{\mathcal{Y}}(\mathbf{u})| = 1$. That is, $\mathbf{y}$ is an edge point if there are points outside $\mathcal{B}$ (i.e., sufficiently far away from $\mathcal{Y}$) that are dominated by $\mathbf{y}$ alone.

In order to determine whether an element $\mathbf{y} \in \mathcal{Y}$ is the sole dominator of some point $\mathbf{u} \in \mathcal{U} \setminus \mathcal{B}$ we observe (see Figure 3) that points which might be dominated by a single $\mathbf{y}_n$ lie "directly above" or "directly to the side" of $\mathbf{y}_n$. With this in mind, candidates for sole domination can be constructed by extending each $\mathbf{y}_n$ along each objective axis in turn into $\mathcal{U} \setminus \mathcal{B}$ and testing how many elements of $\mathcal{Y}$ dominate it. We choose a particular hyper-rectangle to project onto, defined in terms of a point $\mathbf{b}^+ \in \mathcal{U} \setminus \mathcal{B}$, whose coordinates are:

$$b^+_m = b_m + \epsilon, \qquad \epsilon > 0 \tag{5}$$

So long as it is positive, the value of ϵ is immaterial. If the dimension of $\mathbf{y}$ is M, then, as illustrated in Figure 3, M candidate points $\mathbf{u}^m_n$, $m = 1, \ldots, M$ corresponding to $\mathbf{y}_n$ are constructed with coordinates:

$$u^i_n = \begin{cases} y_{ni} & i \neq m \\ b^+_m & i = m \end{cases} \qquad m = 1, \ldots, M \tag{6}$$

By construction $\mathbf{y}_n \preceq \mathbf{u}^m_n$ for all m, but if each of the $\mathbf{u}^m_n$ are dominated by at least one other element of $\mathcal{Y}$ then $\mathbf{y}_n$ is not an edge point. In Figure 3 $\mathbf{u}^1_2$ is dominated by $\mathbf{y}_3$, $\mathbf{y}_4$, $\mathbf{y}_5$ and $\mathbf{u}^2_2$ is dominated by $\mathbf{y}_1$. However, $\mathbf{u}^2_1$ is only (weakly) dominated by $\mathbf{y}_1$, which is therefore an edge point.

Figure 4 presents three examples of edge identification using the attainment surface. Figure 4(a) shows the edge points identified in the concave set shown in Figure 2. Here the method has identified points that agree reasonably with intuition. However on the convex set Figure 4(b) shows that the method fails to identify many of the points that we would hope to lie in the edge set. The reason for this is explained by considering the two edge points labelled α and β. In terms of the first two coordinates y_1 and y_2, α and β are near-optimal. It is therefore unlikely that another element of the set will dominate regions with respect to these coordinates that are not also dominated by α and β, meaning that few of the individuals in the convex population are on the edge; the relatively few individuals that are identified as edge points either have y_1 or y_2 coordinates smaller than those of α or β, or have close to optimal y_3 coordinates. Figure 4(c) shows the edge points of an "hourglass" set, which was constructed from the union of 400 samples of appropriately translated concave and convex sets and thus incorporates concave and convex regions. Here too the edge points of the concave region have been well identified, but points that are intuitively close to the edge of the convex region do not correspond to this definition.

Straightforward application of this definition to high dimensional sets is not successful because almost all elements of the set are identified as edge points. We find that all but two of the 11000 individuals in a mutually non-dominating 9-dimensional criterion set [8, 9] are edge points, with similar outcomes for synthetic sets in many dimensions. This is to be expected, as it is a side effect of the inability of the dominance relation to discriminate between individuals in high-dimensional spaces.

3. FROM CORNERS TO EDGES

We recall Singh *et al.*'s [4] definition of corner points which at first sight appear to be good candidates for extremal points lying on the edge of a mutually non-dominating set. They consider the minimisation of M *functions* $f_i(\mathbf{x})$ where $\mathbf{x}$ is a vector of decision variables. If, when minimising a subset of $k < M$ objectives, there exists a *single* minimising point, then that point is a corner point. Clearly, in non-degenerate cases, there are at least M corner points each corresponding to the minimisation of a single objective, but there are $2^M - 1$ possible combinations to be tested.

The definition of corner points might be adapted to a (finite) *set* of mutually non-dominating points as follows. Let κ denote a set of indices and let $\mathbf{y}^\kappa_n$ be the projection of $\mathbf{y}_n$ onto the indices indicated by κ. If $|\kappa| = k$ then $\mathbf{y}^\kappa_n$ is a k-dimensional vector. Also, let the function nondom($\mathcal{A}$) be the function that returns the maximal set of non-dominated members of the set $\mathcal{A}$:

$$\text{nondom}(\mathcal{A}) = \{\mathbf{u} \in \mathcal{A} \,|\, \nexists(\mathbf{v} \in \mathcal{A} \wedge \mathbf{v} \prec \mathbf{u})\}. \tag{7}$$

Then $\mathbf{y}_i$ is a corner of order k iff $\text{nondom}(\{\mathbf{y}^\kappa_j \,|\, \mathbf{y}_j \in \mathcal{Y}\}) = \mathbf{y}^\kappa_i$; that is, order k corners are the points that dominate all others in at least one of the $\binom{M}{k}$ subsets of k objectives.

Clearly, the extremal points in two-criterion sets (e.g., Figure 1) are also corners. However, with more objectives and a finite set, corners defined like this do not correspond to our intuitive notion of where the corners lie. In fact, for the 3-criterion sets shown in Figure 2 the corners of order 1 are the points that minimise single objectives (marked with circles) and there are no order 2 corners. Note that in the concave case the so-called corners do not lie near to where most people would place the corners.

The reason that there are no order 2 corners is because on projecting onto a pair of criteria there are many non-dominated points; that is $|\,\text{nondom}(\{\mathbf{y}^\kappa_j \,|\, \mathbf{y}_j \in \mathcal{Y}\})| > 1$. Those points which are non-dominated when projected are also candidate edge points. In fact, they turn out to be precisely the same points as those edge points defined in section 2 and are therefore those illustrated in Figure 4.

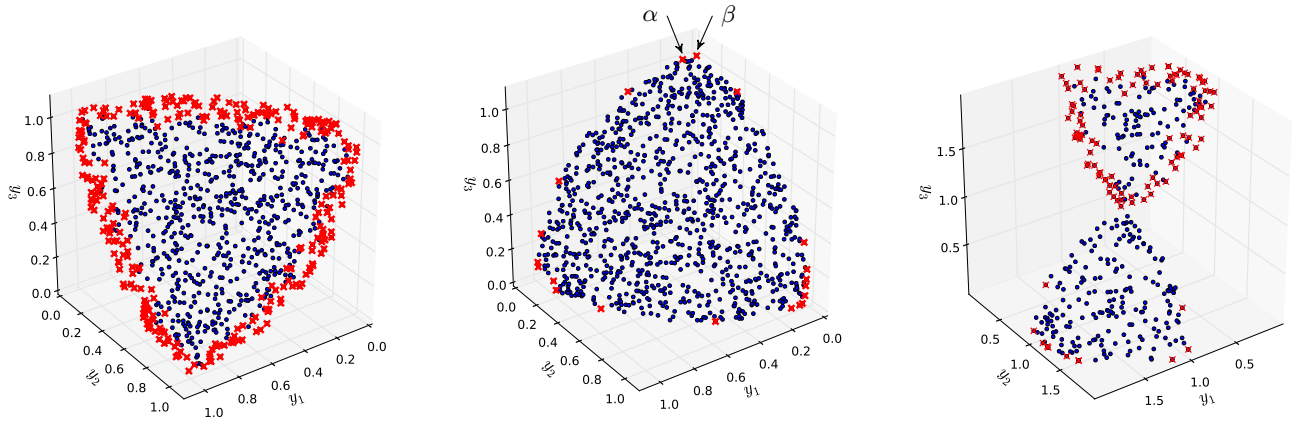

(a) Concave (b) Convex (c) Hourglass

Figure 4: Edge points (crosses) identified by the attainment surface method for concave, convex and hourglass sets.

To see that the non-dominated points following projection onto criterion subsets are the same points as those which extend the range of the attainment surface, suppose that $\mathbf{y}$ is non-dominated when projected onto the indices κ; that is, $\mathbf{y}^\kappa \in \text{nondom}(\{\mathbf{y}_j^\kappa \,|\, \mathbf{y}_j \in \mathcal{Y}\})$. Consequently for criteria $m \in \kappa$, $y_m \leq y'_m$ for all $\mathbf{y}' \in \mathcal{Y}$. Then let $\mathbf{u} \in \mathcal{U} \setminus \mathcal{B}$ be defined as

$$u_m = \begin{cases} y_m & m \in \kappa \\ b_m^+ & m \notin \kappa \end{cases} \qquad m = 1, \ldots, M \tag{8}$$

By construction $\mathbf{y} \prec \mathbf{u}$. On the other hand, if $\mathbf{y}'$ is dominated when projected onto the indicies κ, then there exists a $m \in \kappa$ and $\mathbf{y}'' \in \mathcal{Y}$ such that $y''_m < y'_m$. Therefore if there is a $\mathbf{u} \in \mathcal{U} \setminus \mathcal{B}$ such that $\mathbf{y}' \prec \mathbf{u}$, then $\mathbf{y}'' \prec \mathbf{u}$ so that $\mathbf{y}'$ is not the sole dominator of $\mathbf{u}$. Thus we have shown that if $\mathbf{y}$ is non-dominated when projected onto indices κ then it is the sole dominator of a point $\mathcal{U} \setminus \mathcal{B}$ and is therefore an attainment surface edge point.

Conversely, suppose that $\mathbf{y}$ is an attainment surface edge point. Then there exists a $\mathbf{u} \in \mathcal{U} \setminus \mathcal{B}$ such that $\mathbf{y} \prec \mathbf{u}$, but no other $\mathbf{y}' \in \mathcal{Y}$ dominates $\mathbf{u}$. Let κ be the indices for which

$$y_k \leq u_k < y'_k \tag{9}$$

Clearly κ is not empty because $\mathbf{y} \prec \mathbf{u}$ and $\mathbf{y}' \not\prec \mathbf{u}$ for all $\mathbf{y}' \in \mathcal{Y}$ ($\mathbf{y}' \neq \mathbf{y}$). Then (9) shows that when projecting onto the criteria κ, we have $\mathbf{y}^\kappa \prec (\mathbf{y}')^\kappa$ for all $\mathbf{y}' \in \mathcal{Y}$, which establishes that attainment surface edge points are non-dominated when projected onto some criterion subset.

We anticipate that the equivalence we have shown between points that extend the attainment surface and points which are non-dominated when projected onto some subset of the criteria will be of use in evolutionary multi-objective algorithms, such as [5, 6, 7], which seek to preserve diversity and the spread of the estimated Pareto front by preferentially retaining and perturbing solutions on the periphery of the solution set.

In section 5 we return to the notion of identifying edges through projections onto criterion subsets, but we first consider an alternative way of identifying edges.

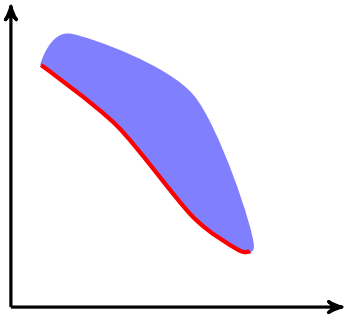

Figure 5: Non-dominated points identify the boundary of a set. Non-dominated regions of the set in the plane are marked in red.

4. DOMINANCE-BASED EDGES WITH ROTATIONS

The motivation for this method comes from the fact that, as illustrated in Figure 5, the non-dominated elements of a set define a portion of the set's edge. The edges of the set could therefore be identified as the union of the non-dominated points as the set is rotated.

In general the points of a multi-criterion non-dominated set are not co-planar, so in order to use this idea with many-criterion non-dominated sets we first project the points in $\mathcal{Y}$ onto a plane in which they remain mutually non-dominating.

Without loss of generality, we assume that the elements of $\mathcal{Y}$ are non-negative ($y_{nm} \geq 0$ for all n, m). Then the *simplex* defined by the numbers $\{\lambda_m : \lambda_m > 0\}_{m=1}^M$ is the portion of the (hyper-) plane which lies in the positive orthant and which intersects the coordinate axes at distances λ_m from the origin, as illustrated in Figure 6. The simplex is therefore the segment of the plane in the positive orthant defined by

$$\mathbf{n} \cdot \mathbf{y} = d \qquad y_m \geq 0, \; m = 1, \ldots, M \tag{10}$$

where the elements of the unit vector $\mathbf{n}$ normal to the simplex are $n_m = d/\lambda_m$ and the perpendicular distance to the

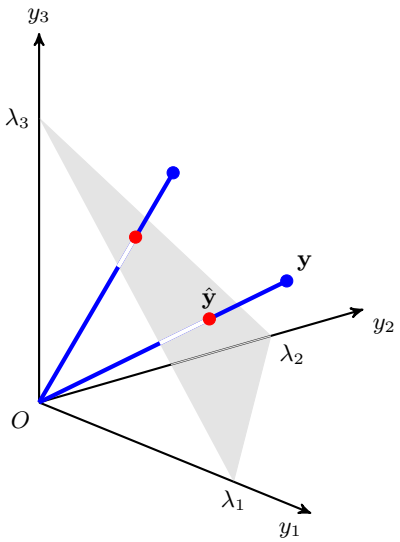

Figure 6: Projection of a point $\mathbf{y}$ onto the simplex defined by $\{\lambda_1, \lambda_2, \lambda_3\}$; $\mathbf{y}$ is projected to $\hat{\mathbf{y}}$.

origin d can be found as

$$d^{-2} = \sum_{m=1}^{M} \lambda_m^{-2}. \tag{11}$$

Points are then projected onto the simplex by:

$$\hat{\mathbf{y}}_n = \mathbf{y}_n / (\mathbf{y}_n \cdot \mathbf{n}), \tag{12}$$

and we emphasise that the projection of $\mathcal{Y}$ results in a set of mutually non-dominating points. The choice of the λ_m clearly affects the particular projection, but since the dominance relations between the points are unaffected, the method is not sensitive to the precise values. We have used $\lambda_m = \text{median}_{\mathbf{y} \in \mathcal{Y}}\, y_m$ here.

To identify the edges of this set, we now rotate the $\hat{\mathbf{y}}_n$ in the plane of the simplex and identify the non-dominated points in each rotation.

Coordinates for the $\hat{\mathbf{y}}_n$ on the simplex may be simply found by projecting the $\hat{\mathbf{y}}_n$ onto their principal components or equivalently by singular value decomposition. Let $\hat{\mathbf{Y}} = [\hat{\mathbf{y}}_1 - \boldsymbol{\mu}, \ldots, \hat{\mathbf{y}}_N - \boldsymbol{\mu}]$ be the matrix whose columns are the mean-centred $\hat{\mathbf{y}}_n$ and let $\hat{\mathbf{Y}} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T$ be the singular value decomposition of $\hat{\mathbf{Y}}$. Then the first $M-1$ columns of the orthonormal matrix $\mathbf{U}$ span the subspace of the simplex (the last column is a vector normal to the simplex). Coordinates of $\mathbf{y}_n$ mapped to the simplex are thus

$$\tilde{\mathbf{y}}_n = \mathbf{U}_{M-1}^T (\hat{\mathbf{y}}_n - \boldsymbol{\mu}) \tag{13}$$

where $\mathbf{U}_{M-1}$ denotes the matrix of the first $M-1$ columns of $\mathbf{U}$.

Rather than exhaustively quartering all rotations, we generate rotations $\mathbf{Q}$ at random. Uniformly distributed random rotations are generated by, for example, a QR decomposition of $(M-1) \times (M-1)$ dimensional matrices whose elements are Gaussian distributed. The signs of the columns of $\mathbf{Q}$ are arbitrary, so in all 2^{M-1} rotations can be cheaply generated from a single QR decomposition by appropriately flipping the signs of the columns of $\mathbf{Q}$.

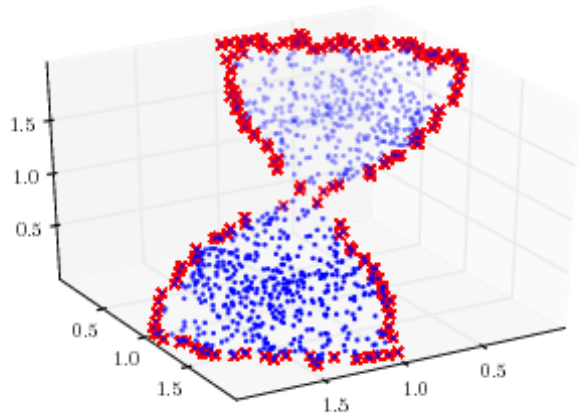

Figure 7: Edge points of the hourglass set identified from non-dominated points after projection onto the simplex and 64 random rotations and associated flips.

As each rotation is generated the non-dominated points $\text{nondom}(\{\mathbf{Q}\tilde{\mathbf{y}}_n \,|\, \mathbf{y}_n \in \mathcal{Y}\})$ are found and the corresponding $\mathbf{y}$ added to the identified edge points.

This method of determining edges is demonstrated for the hourglass set in Figure 7. It is clear that a reasonable set of points has been identified as edge points, both on the upper, concave portion and on the lower, convex section. However, the relationship of these edge points to those found with the attainment surface or by projection onto subsets of the criteria is not clear. We note also that if $\mathcal{Y}$, when projected onto the simplex, has "deep and narrow" concavities then this method will be unable to locate edge points in the concavities.

We also applied this method to the task of identifying the edge of the 9-criterion radar data. However, like the other methods, it too identifies all points as being on the edge. Again, this is due to the relative lack of discrimination provided by the dominance relation in high dimensional spaces.

5. CRITERION SUBSET EDGES

In section 3 we identified edges as the non-dominated points after projection onto k of the M criteria. As we showed there, these points are the points which extend the range of the attainment surface. However, reference to Figure 2 for example shows that points which *maximise* one of the criteria also lie on the intuitive edges of concave and convex fronts. In this section we therefore extend the criterion for a point to be an edge point by including points that, after projection onto k of the criteria, are not dominated if the criteria were to be *maximised* rather than minimised.

To be specific, let the function $\overline{\text{nondom}}(\mathcal{Y})$ be the function which returns the maximal set of non-dominated members of $\mathcal{Y}$ under maximisation:

$$\overline{\text{nondom}}(\mathcal{Y}) = \{\mathbf{y} \in \mathcal{Y} \,|\, \nexists(\mathbf{v} \in \mathcal{Y} \wedge \mathbf{v} \succ \mathbf{y})\} \tag{14}$$

Then, with $\mathbf{y}^\kappa$ denoting the projection of $\mathbf{y}$ onto the criteria specified by κ, our final definition of the set of edge points

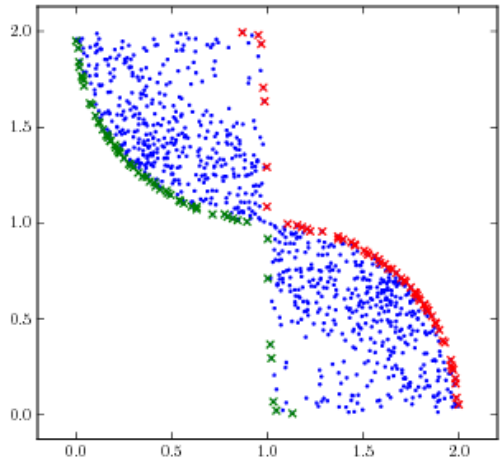

Figure 8: The hourglass data projected onto criteria y_2 and y_3. Points which are non-dominated under minimisation and maximisation are marked with green and red crosses respectively.

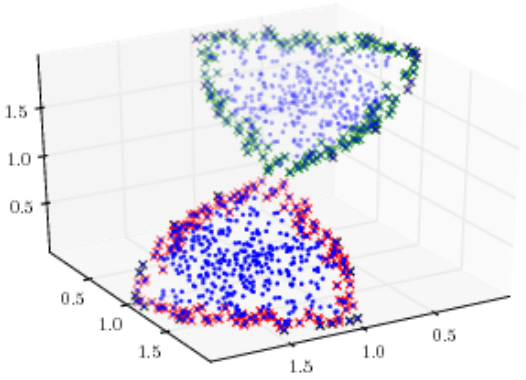

Figure 9: Edge points of the hourglass data identified as non-dominated under minimisation or maximisation after projection onto criterion subsets. Points which are non-dominated under minimisation and maximisation are marked with green and red crosses respectively; those which are non-dominated under both are marked in black.

order $k = |\kappa|$ is:

$$\{\mathbf{y}_n \,|\, \mathbf{y}_n^\kappa \in \text{nondom}(\{\mathbf{y}_n^\kappa \,|\, \mathbf{y}_n \in \mathcal{Y}\})\} \bigcup \{\mathbf{y}_n \,|\, \mathbf{y}_n^\kappa \in \overline{\text{nondom}}(\{\mathbf{y}_n^\kappa \,|\, \mathbf{y}_n \in \mathcal{Y}\})\}. \quad (15)$$

As an illustration, Figure 8 shows the projection of the hourglass data onto criteria y_2 and y_3. Points which are non-dominated under minimisation and maximisation are marked with green and red crosses respectively and those that are non-dominated under both are marked in black. It can be seen that these correspond to different regions of extremal points. Figure 9 shows all the edge points identified after projections onto all the criterion subsets. As the figure illustrates, this definition of edge points has identified a uniform spread of points corresponding to what we intuitively identify as edges.

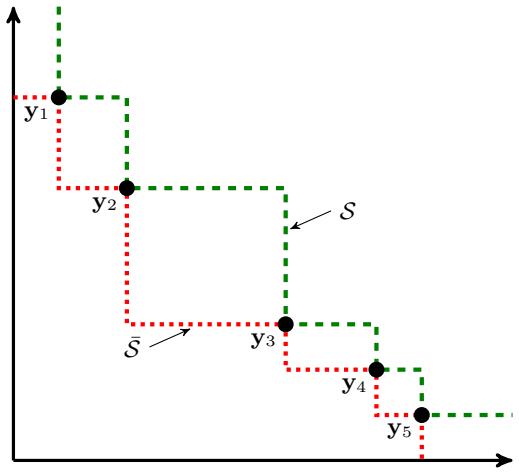

Figure 10: Attainment surfaces under minimisation and maximisation. The attainment surfaces of a mutually non-dominating set of points $\mathcal{Y} = \{\mathbf{y}_i\}$ under minimisation and maximisation are shown by green dashed and red dotted lines respectively.

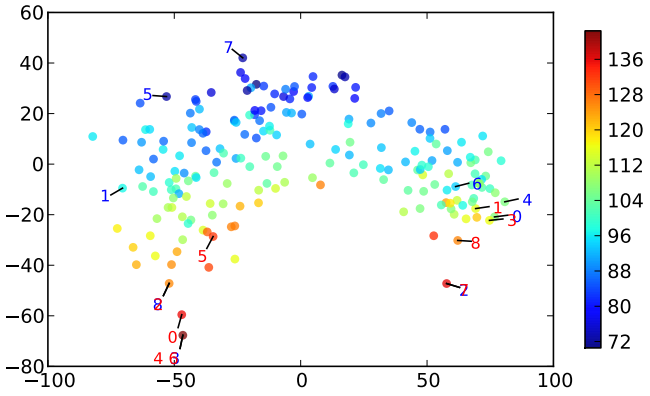

Figure 11: Dominance distance MDS visualisations of 200 samples from the radar data. Solutions are coloured by their average rank. The best and worst solutions on each objective are labelled by that objective in blue and red respectively.

The attainment surface of $\mathcal{Y}$ under maximisation $\bar{\mathcal{S}}$ may be defined as the boundary of the region dominated under maximisation by the elements of $\mathcal{Y}$. Figure 10 illustrates the attainment surfaces under minimisation and maximisation for a set of mutually non-dominating points in the plane. Whereas the attainment surface under minimisation is a conservative interpolation of the set [7] in the sense that every point in $\mathcal{S}$ is weakly dominated by an element of $\mathcal{Y}$, the attainment surface under maximisation can be seen to be an optimistic interpolation of $\mathcal{Y}$. Using the same arguments as presented in section 3, the edge points identified as non-dominated under maximisation are exactly those which extend the range of the attainment surface under maximisation. We note that for two criteria, the extremal points are edge points because they extend both $\mathcal{S}$ and $\bar{\mathcal{S}}$; of course, this is not the case with more than two criteria.

To visualise the edges of the radar dataset we first project it into two dimensions. Rather than use the objective values directly, we seek to capture the order relations between

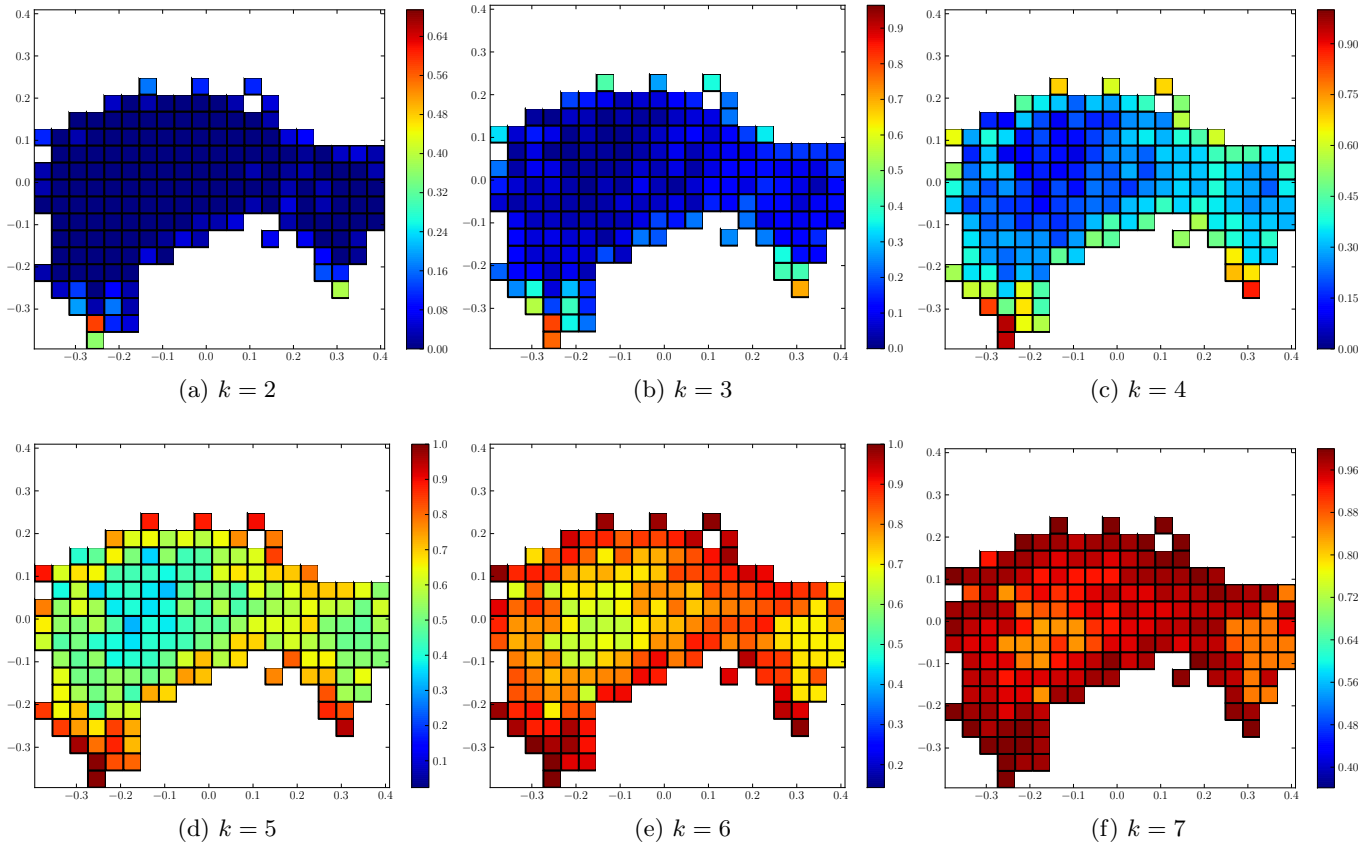

(a) $k = 2$ (b) $k = 3$ (c) $k = 4$

(d) $k = 5$ (e) $k = 6$ (f) $k = 7$

Figure 12: Cellular visualisations of the radar population embedded into 2 dimensions with dominance distance MDS. Cells are coloured according to the median frequency with which edge points in the cell are edge points in a k-criterion subset.

the elements of the set. This is done by defining a *dominance distance* between elements, which quantifies the extent to which two elements are on average greater than, less than or equal to other elements in the set. The dominance distance is formally a metric [3]. Metric multi-dimensional scaling (MDS, [11, 12]) is then used to find an embedding in Euclidean space which preserves the dominance distances. Figure 11 shows the dominance distance MDS visualisation of 200 solutions of the radar data, in which solutions are coloured by their average rank. As the figure shows, the visualisation provides a topographic low-dimensional representation of the data. Additional details and visualisations are given in [3].

In common with the other definitions of edge points, with many criteria the vast majority of points are edge points. Table 1 shows the variation of the number of edge points identified in 2000 radar data points. Clearly the fraction of edge points grows with the number of criteria. For comparison the table also shows the number of order k edge points for 2000 points uniformly distributed on the positive orthant of a spherical shell in 9 dimensions; here the convex nature of the front is apparent in the disproportionate number of edge points from maximisation.

Figure 12 shows the dominance distance embedding of the radar data divided into cells, each of which is coloured according to the median frequency of edge cells for orders $k = 2, \ldots, M - 1 = 8$. As shown, identified edge points tend to occur more frequently close to the edges of the dominance distance MDS projection. The projection and this definition of edge points therefore concur about where the boundaries of the set lie. For all orders the maximum density of edge points is found on the horns of the crescent, which is also where the points maximising or minimising individual criteria are mostly located, as shown in Figure 11.

An additional interesting feature revealed by this visualisation of the edge points is the region of higher edge point

	Radar front			Convex front		
$\lvert\kappa\rvert$	Min	Max	Both	Min	Max	Both
2	469	201	87	151	372	49
3	1611	1182	951	498	1691	482
4	1949	1793	1745	941	2000	941
5	1993	1950	1943	1693	2000	1693
6	1998	1986	1984	1693	2000	1693
7	1998	1996	1994	1894	2000	1894
8	1998	1998	1996	1981	2000	1981

Table 1: The number of order $k = |\kappa|$ edge points in 2000 solutions from the radar estimated Pareto front and from a 9-criterion convex mutually non-dominating set. Columns show the number of edge points arising from non-dominance under minimisation, under maximisation and the number that are non-dominated under both.

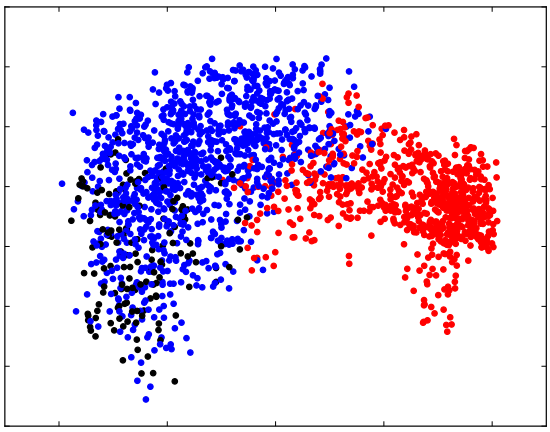

Figure 13: Dominance distance MDS embedding of the radar data with points coloured according to the type of objective for which they are most highly ranked. Red: range. Blue: velocity. Black: duration. There is a distinct boundary between solutions which best optimise range criteria and those which best optimise velocity criteria or transmission time. This boundary is in the same position as the high density of edge points shown in Figure 12.

density crossing the embedding from top to bottom towards the righthand third of the embedding. The criteria on which the radar data is optimised divide into three groups. One group (f_1, f_3, f_5 and f_7) relate to the range at which a radar can discern targets. The second group (f_2, f_4, f_6 and f_8) relates to the velocity at which the radar can be moving and still discern targets, while the final group (f_9) comprises a single objective, the transmission time of the radar waveform. Individual solutions tend to optimise one of these groups of criteria at the expense of the others [3]. Plotting the solutions coloured according to which of these groups is best optimised, Figure 13, shows that the high edge point density region corresponds to the transition between solutions that best optimise the range objectives and those that best optimise the velocity objectives.

6. CONCLUSION

Three of our definitions of edge points are closely related to the attainment surface, which plays a fundamental role in understanding which points lie close to the edges of a mutually non-dominating set. We have shown that the definition of edge points as those points which extend the range of the attainment surface yields precisely the points which are non-dominated in a projection onto a subset of the criteria.

Note that the edge points identified as the maximal mutually non-dominating points after projection onto subsets of the criteria have also been considered by di Pierro *et al.* in their k-preference ordering [13], in which a many-objective point is considered important if it remains non-dominated when projected onto criterion subsets. Although somewhat counter-intuitive we have shown that points which remain non-dominated under *maximisation* after projection onto criterion subsets are also important, because they too are on the edges of the set. They play a particularly important role in convex-shaped sets.

Of the four candidate definitions of edge points we prefer the final one (§5), which identifies edge points as those points which, after projection onto k criteria, are not dominated under either minimisation or maximisation. As the examples show, it most fully captures the intuitive notion of edges and is simple to compute.

As the number of criteria increases and almost all points are edge points and thus extend the range of the attainment surface, we emphasise that regions with a high frequency of edge points correspond to the boundaries of low-dimensional visualisations, and furthermore reveal previously unknown structure in a many-objective optimisation problem.

7. REFERENCES

[1] Ishibuchi, H., Tsukamoto, N., Nojima, Y.: Evolutionary Many-objective Optimization: A Short Review. In: Proc. IEEE Congress on Evolutionary Computation. (2008) 2419–2426

[2] Kudo, F., Yoshikawa, T.: Knowledge Extraction in Multi-objective Optimization Problem based on Visualization of Pareto Solutions. In: Proc. IEEE Congress on Evol. Comp. (2012) 860–865

[3] Walker, D., Everson, R., Fieldsend, J.: Visualising mutually non-dominating solution sets in many-objective optimisation. IEEE T-EVC (2012) doi: 10.1109/TEVC.2012.2225064.

[4] Singh, H., Isaacs, A., Ray, T.: A Pareto corner search evolutionary algorithm and dimensionality reduction in many-objective optimization problems. IEEE T-EVC **15**(4) (2011) 539–556

[5] Zitzler, E., Laumanns, M., Thiele, L.: SPEA2: Improving the Strength Pareto Evolutionary Algorithm for Multiobjective Optimization. In: Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001). (2002) 95–100

[6] Knowles, J.D., Corne, D.W.: Approximating the Nondominated Front Using the Pareto Archived Evolution Strategy. Evolutionary Computation **8**(2) (2000) 149–172

[7] Smith, K., Everson, R., Fieldsend, J., Misra, R., Murphy, C.: Dominance-based multi-objective simulated annealing. IEEE T-EVC **12**(3) (2008) 323–343

[8] Hughes, E.J.: Radar waveform optimisation as a many-objective application benchmark. In: Proc 4th International Conference on Evolutionary Multi-criterion Optimization. (2007) 700–714

[9] Hughes, E.J.: Best Known Non-dominated Results of Radar Waveform Optimisation. `code.evanhughes.org` (2007)

[10] Zitzler, E.: Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications. PhD thesis, ETH Zurich (1999)

[11] Sammon, J.W.: A Nonlinear Mapping for Data Structure Analysis. IEEE Transactions on Computers **18**(5) (1969) 401–409

[12] Webb, A.R.: Statistical Pattern Recognition, 2nd Edition. John Wiley & Sons (October 2002)

[13] di Pierro, F., Khu, S.T., Savic, D.: An Investigation on Preference Ordering Ranking Scheme in Multiobjective Evolutionary Optimization. IEEE T-EVC **11**(1) (2007) 17–45