Rotated Test Problems for Assessing the Performance of Multi-objective Optimization Algorithms

Antony W. Iorio
School of Computer Science and I.T., RMIT University
Melbourne VIC 3000 Australia
iantony@cs.rmit.edu.au

Xiaodong Li
School of Computer Science and I.T., RMIT University
Melbourne VIC 3000 Australia
xiaodong@cs.rmit.edu.au

ABSTRACT

This paper presents four rotatable multi-objective test problems that are designed for testing EMO (Evolutionary Multi-objective Optimization) algorithms on their ability in dealing with parameter interactions. Such problems can be solved efficiently only through simultaneous improvements to each decision variable. Evaluation of EMO algorithms with respect to this class of problem has relevance to real-world problems, which are seldom separable. However, many EMO test problems do not have this characteristic. The proposed set of test problems in this paper is intended to address this important requirement. The design principles of these test problems and a description of each new test problem are presented. Experimental results on these problems using a Differential Evolution Multi-objective Optimization algorithm are presented and contrasted with the Non-dominated Sorting Genetic Algorithm II (NSGA-II).

Categories and Subject Descriptors

G.1.6 [**Mathematics of Computing**]: Numerical Analysis—*Optimization*

General Terms

Experimentation, Performance, Measurement

Keywords

Parameter interactions, Multi-objective Optimization

1. INTRODUCTION

Traditionally, multi-objective problems have been solved using classical techniques, but such approaches typically have many limitations with respect to the types of problems they can solve, and may not even be able to find optimal solutions for problems which have no known functional representation. A relatively recent approach to solving multi-objective problems has been the application of population based Evolutionary Algorithms (EAs). There is a natural synergy between a population based evolutionary approach, which searches a number of points in the decision space simultaneously, and a multi-objective problem, which has a Pareto-optimal set of solutions. Evolutionary approaches can provide a set of solutions which closely approximates the Pareto-optimal set. Despite their apparent successes, there is a disparity between how we evaluate EMO algorithms with test problems with specified properties, and the characteristics of many real-world problems.

GECCO'06, July 8–12, 2006, Seattle, Washington, USA.

1.1 Test problems for EMO Algorithms

Typically, in order to assess the performance of EMO algorithms in a variety of fitness landscapes, test problems with a variety of characteristics are employed. For instance, one may wish to evaluate the performance of an algorithm with respect to the concavity of the Pareto-optimal front, or one may employ test problems with discontinuous Pareto-optimal fronts, sparsely distributed solutions near the Pareto-optimal front, or non-uniform mappings between the decision and objective space. The ZDT series of test problems [13], built on the framework proposed in [1], exhibit some of these characteristics. In [1] the problem features that may cause difficulty for an EMO algorithm are discussed, and a framework is proposed for construction of test problems exhibiting such features. Recently, there has been further progress in the construction of test problems for multi-objective optimization. Although many approaches have focussed on the construction of only two-objective problems, other approaches deal with the construction of test problems with more than two objectives in [6] and [3]. A framework is also proposed in [9] for constructing arbitrary user specified Pareto-optimal sets in the decision space, which map to Pareto-optimal fronts in the objective space. In the area of multi-objective combinatorial optimization, instance generators have been proposed for the quadratic assignment problem [8]. These generators are useful for studying parameter interactions in combinatorial problems.

1.2 Problems with parameter interactions

Before discussing the previous work in the area of test problems with parameter interactions, it is necessary to define what we mean by a linearly separable problem. Parameter interactions occur in a problem because the parameters in the problem are not linearly separable. The geometric interpretation of what constitutes a linearly separable problem is where the n parameters of a problem in n-dimensional space can be separated by an $n-1$ dimensional hyperplane. We will see in Section 2 how linear separability occurs with a simple ellipsoid function, and how rotation can make the function non-separable.

Previous work in [12] demonstrated the importance of evaluating single objective Evolutionary Algorithms on rotated problems in order to test their rotationally invariant behaviour, the contention being that evaluations of EAs should be independent of any coordinate system. Furthermore, in the EMO domain, NSGA-II was demonstrated to perform poorly on a simple rotated problem [2] and the importance of parameter interactions in multi-objective problems was discussed in [7]. If EMO algorithms can fail on such a simple problem, but succeed when the problem is aligned with the principle coordinate system of the decision space, then obviously many reported results on such test problems are potentially misleading. In [12] a strong case is presented for all EA evaluations to be independent of a particular coordinate system, and our contention is that this should equally be true for the evaluation of EMO algorithms.

Recently, the issue of parameter interactions in multi-objective problems has garnered more interest; in [3] the non-separability of multi-objective problems is demonstrated to be an important characteristic lacking in existing multi-objective test suites, and a method is described for constructing problems with parameter interactions for an arbitrary number of objectives. In [11] a differential evolution multi-objective algorithm was evaluated on a number of the proposed non-separable and separable problems proposed in [3].

1.3 Overview

Each of the real-coded ZDT test problems has a Pareto-optimal set in the decision space which is aligned with the principle coordinate system. This makes the problem easier, because the objectives can be solved in stages. The purpose of this paper is to propose a complementary series of test problems to the ZDT series of problems, using the approach described in [1], and to provide some further insight into rotated multi-objective problems. The primary motivation for this is that most real world problems have parameter interactions, and ideally we would also like a variety of test problems which have parameter interactions as well. The framework employed for constructing the ZDT problems is commonly used by practitioners, and it is easily employed in the construction of test problems.

The proposed problems in this paper can be arbitrarily rotated on any axis of the decision space. Through a rotation of the coordinate system, parameter interactions can be introduced to the problem. When the proposed problems are rotated in the decision space, the Pareto-optimal set is rotated accordingly. In Section 2.1 we will see that when any solution in the set is perturbed independently with respect to any of the decision variables, it can only be perturbed to non-dominated solutions which are not Pareto-optimal.

The proposed series of rotated problems presented in this paper will provide a means of assessing the performance of EMO algorithms which is not biased towards a particular orientation of the problem with respect to the coordinate axes. This provides EMO practitioners with a reliable indication of the performance of their algorithm irrespective of the orientation of the problem in the decision space.

In the following section we will describe in more detail how rotation makes single objective and multi-objective problems harder to optimize. In Section 3 we describe the proposed problem suite, and how the problems were constructed. For a comparison with the NSGA-II algorithm, a new Differential Evolution approach to EMO is briefly introduced in Section 4 along with the experiments conducted. This is followed by a discussion of the results and the utility of the proposed problem suite in Section 5 and concluding remarks in Section 6.

2. ROTATED PROBLEMS

Parameter interactions can be introduced in other types of problems, such as combinatorial problems. For the purposes of this study, we are only considering multi-objective problems with real-valued decision variables, where one or more objectives have non-linear parameter interactions.

Before elucidating upon rotated multi-objective problems, we will first consider the effect of rotation on a simple single objective ellipsoid minimization problem defined in Equation (1). Figure 1(a) also presents a contour plot of the function.

$$f(x_1, x_2) = x_1^2 + a_0 x_2^2 \quad (1)$$

$$f(x_1, x_2) = x_1^2 + a_1 x_1 x_2 + a_0 x_2^2 \quad (2)$$

The ellipsoid problem in Equation (1) has a global minimum located at $x_1 = 0$ and $x_2 = 0$, at the origin O, of the principle coordinate axes. It is apparent from the contour plot in Figure 1(a), that this function is aligned with the principle coordinate axes, and is linearly separable with respect to the two decision variables. The two components, x_1^2 and $a_0 x_2^2$, of Equation (1), can be solved as independent minimization problems. A search algorithm only needs to perturb the variables x_1 and x_2 independently in order to find the global optimum for this problem. If the ellipsoid function from Equation (1) is rotated away from the principal coordinate axes (Figure 1(b)), the decision variables become non-separable through the introduction of parameter interactions in Equation (2). Parameter interactions are introduced through the term $a_1 x_1 x_2$. Progress towards the global optimum can only proceed efficiently by making simultaneous improvements with respect to all parameter values.

In Figure 1, the contour represents a region of constant fitness. The point A in Figure 1(a) can be perturbed along the x_1 and x_2 axes, and any location along the dashed lines will be an improvement over any point along the contour. In Figure 1(b) it is apparent that progress from perturbing the rotated point A' will be lower. This is because the interval of potential improvement for each of the decision variables is reduced, and as a result, the progress of the search through independent perturbations will be reduced.

Furthermore, the search can easily be trapped on the line segment that bisects the ellipsoid lengthwise, as can be seen in Figure 2. Any point on this line segment can only move to another point which evaluates to a better solution, by making simultaneous improvements on each decision variable. For example, in Figure 2(a), the point A can be independently perturbed on the x_1 axis to find the global minimum located at the origin of the coordinate system. The same point A', in Figure 2(b), after rotation cannot progress to a point of improved fitness by only moving along the direction

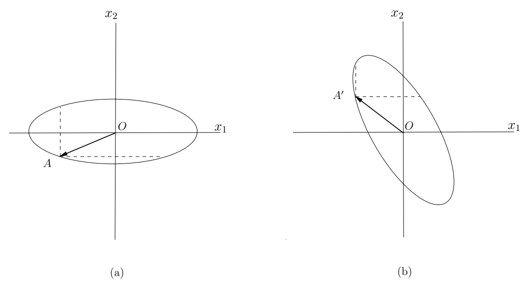

Figure 1: The contour plots demonstrate how rotation can reduce the interval of possible improvement represented by the dashed line segments. Point A in (a) and A' in (b) represent the same point on the contour plot before and after rotation respectively. The contour represents a region of constant fitness.

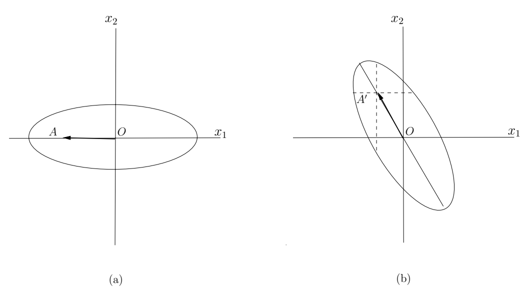

Figure 2: The contour plots demonstrate how rotation can trap points along the line segment bisecting the ellipsoid. Point A in (a) and A' in (b) represent the same point on the contour plot before and after rotation respectively.

of the principle coordinate axes because any such perturbation will be to a point of lower fitness in the objective space. Typically the basin of attraction can be found comparatively easily, but the search can become trapped when the problem is non-separable. Only a simultanenous improvement in all parameters will result in the discovery of fitter solutions in this situation. On these types of problems, the small mutation rates frequently used in Genetic Algorithms are known to be even less efficient than a random search [12]. Evolutionary Strategies have been relatively successful at solving these types of problems, but they require the learning of appropriate correlated mutation step sizes and it can be rather computationally expensive when the decision space dimension becomes large [10].

2.1 Rotated Multi-objective Problems

Although one might intuitively expect that in the multi-objective domain the situation encountered is the same, the situation is not quite as simple. In order to highlight the difference between rotated single-objective and multi-objective problems, we will consider a simple multi-objective problem. This will also facilitate our understanding of the effect of rotation on multi-objective problems where non-dominated solution sets are sought by a search algorithm. The situation is analogous to the single objective domain, where a search algorithm with independent perturbations on each decision variable will have trouble finding more optimal solutions.

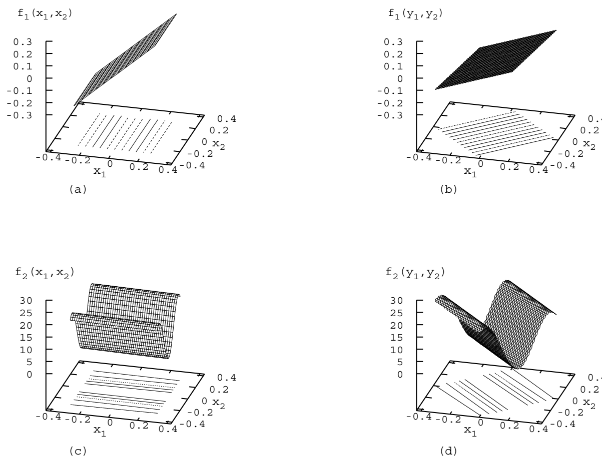

Figure 3: The effect of a 45-degree rotation on the x_1x_2 plane of problem R1. Before rotation, the functions are aligned with the coordinate system ((a) and (c)), and after rotation they are not ((b) and (d)).

Figure 3 shows a simple bi-objective optimization problem with a 2-dimensional decision space. This is problem R1, defined in Section 3.2.

The problem is characterised by a slightly inclined trough in objective f_2. Objective f_1 is a plane with a gradient sloping in an opposing direction to the incline of objective f_2. The Pareto-optimal set is represented by a line segment bisecting the decision space in objective f_2 and f_1 respectively. The decision space is subject to a rotation matrix R.

Consider the contour plot of a non-rotated version of this problem in Figure 4, where a point, O, is a member of the Pareto-optimal set. If the point O is perturbed in the direction of $\vec{OA}$, it is towards points which evaluate lower with respect to objective f_1 (Figure 4(a)) and higher with respect to objective f_2 (Figure 4(b)). If it is perturbed in the direction of $\vec{OD}$, it will be towards points which evaluate higher with respect to objective f_1 and lower with respect to objective f_2. Such perturbations are with respect to the parameter x_1 only, and this is the only such perturbation required in order to discover other Pareto-optimal solutions which are located on the line segment bisecting the contour plots for objectives f_1 and f_2. It is apparent that such a Pareto-optimal solution can easily be perturbed towards other Pareto-optimal solutions when the problem is aligned with the principle coordinate axes. However, after the problem has been rotated, it becomes more difficult to find Pareto-optimal solutions through independent perturbations of individuals. Consider Figure 5, where point O' represents the point O after rotation. Through independent perturbations of decision space parameters, the point O can perturb to other non-dominated solutions in the direction of $\vec{O'A'}$, $\vec{O'B'}$, $\vec{O'C'}$, and $\vec{O'D'}$. A perturbation in the direction of $\vec{O'A'}$ leads to points which evaluate lower on objective f_1, but higher on objective f_2. This is similarly true for perturbations in the direction of $\vec{O'B'}$. A perturbation in the direction of $\vec{O'C'}$ leads to points which evaluate higher on objective f_1, but lower on objective f_2. This is

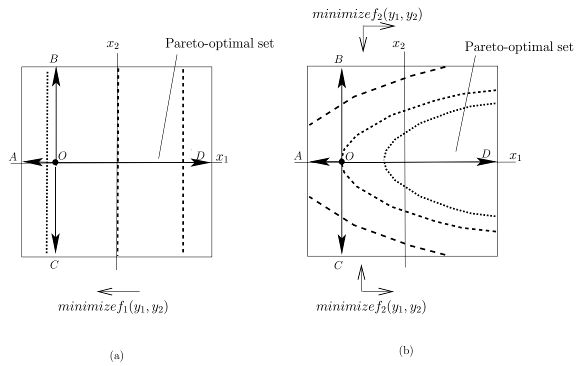

Figure 4: The contour plot of non-rotated R1 in (a) represents function f_1, and the contour plot in (b) represents objective function f_2. The dashed lines represent regions of constant value with respect to the objective function evaluation. Smaller dash sizes represent lower evaluations on the objective functions. The point O, in the Pareto-optimal set, can be perturbed to other Pareto-optimal points by perturbing the decision variable x_1.

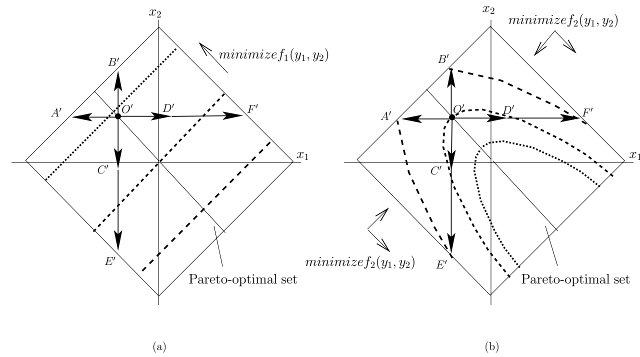

Figure 5: The contour plot of rotated R1 in (a) represents function f_1, and the contour plot in (b) represents objective function f_2. After rotation, the point O', in the Pareto-optimal set, can be perturbed to other Pareto-optimal points by perturbing the decision variables x_1 and x_2 simultaneously along the line segment which bisects both feasible regions.

also true for perturbations in the direction of $\vec{O'D'}$. Unfortunately each of these perturbations leads to non-dominated solutions which skew away from the Pareto-optimal set. The situation becomes even worse if the perturbation extends to $\vec{C'E'}$ or $\vec{C'F'}$, because individuals in this region evaluate higher with respect to objective f_1 and objective f_2, and are dominated by the point at O', as a result. In actuality, the problem in Figure 3 that this analysis is based on, has an extremely small region relative to the feasible space, where non-dominated solutions can be located in the direction of $\vec{O'C'}$ and $\vec{O'D'}$. Non-dominated solutions in these directions only becomes increasingly likely as the orientation of the problem approaches alignment with the principle coordinate axes. As the orientation of the Pareto-optimal front aligns with the axis x_1, the line vector $\vec{O'C'}$ extends further and more non-dominated solutions can be discovered in this region more easily. Secondly, such non-dominated solutions will be close to the Pareto-optimal set. This is similarly true for the line vector $\vec{O'D'}$, as the Pareto-optimal front aligns with the axis x_2. In other words, a rotation of the problem which results in the Pareto-optimal set not being aligned with any principle coordinate axis, makes it difficult to discover other Pareto-optimal solutions when only independent perturbations of decision variables can occur.

In the presence of only independent perturbations, there is a tendency for points to be discovered in the direction of lower f_1 evaluations, and higher f_2 evaluations, pushing the non-dominated solution set away from the Pareto-optimal set and degrading the search over time. As a result of this behaviour, the search can become trapped in the Pareto-optimal region and fail to find more non-dominated solutions in the Pareto-optimal set. Progress in covering the Pareto-optimal front becomes extremely slow. This effect was also apparent in [4] and [2], where the NSGA-II produced poor coverage of the Pareto-optimal front on the rotated uni-modal multi-objective problem. Any multi-objective optimization algorithm which is not rotationally invariant, will exhibit such behaviour.

3. CONSTRUCTION OF ROTATED TEST PROBLEMS

Each of the test problems proposed was designed using the frame work proposed in [1]. The g function is responsible for affecting convergence to the Pareto-optimal front. The h function specifies the shape of the Pareto-optimal front, whether the problem is rotated or not. The Pareto-optimal set can be determined by setting the g function to 1.0, and evaluating f_1 over the range of feasible solutions. Diversity is also affected by the f_1 function.

In order to construct a problem with parameter interactions, the problem must have at least one non-linear function. With this consideration, the rotatable test problems we have proposed in this paper will have at least one non-linear function in at least one of the objective functions. Each of the test problems also sets f_1 and f_2 to large values if f_1 is outside the ranges specified in the problem descriptions. This is important because rotation may push a function evaluation outside the range desired by the experimenter [1].

In order to construct a rotated problem, one must also be careful that under rotation the problem can still evaluate to a meaningful result. One should avoid situations where a rotation transformation results in decision variables which take negative values, and are then subjected to a square root function for instance. This can be achieved by avoiding functions which would result in such a situation, or by offsetting the variable value within the function so a negative value never results. It should also be noted that the methodology for performing a rotation can be applied to other problems from the literature, and is not limited to the ZDT problems. The only considerations that need to be addressed are related to making sure the evaluation of the rotated vector yields a result that can still be evaluated. The Pareto-optimal set can still be determined by the means specified within the test problem construction methodology employed by the user.

3.1 Generating uniform random rotations

In order to achieve a completely unbiased assessment of an algorithm on a problem which is rotated, one must guarantee a uniformly distributed random rotation. Algorithm 1 outlines the procedure for generating a random orthnormal basis which is used to introduce parameter interdependencies into a problem by rotating the parameter vector. This is a transformation which does not change the fitness landscape of the problem domain because it is an *isometry* transformation; it preserves distances between points, and as a result it also preserves angles. A uniform distribution of points on the surface of a hypersphere are possible when the orthonormal basis $R_1, ..., R_m \in \mathbb{R}^m$, is used to rotate a point in m-dimensional space. This technique also makes it possible to randomly and uniformly rotate a decision space vector $\vec{x}$, by using matrix multiplication $R\vec{x}$, so that there is no bias for any particular coordinate axis. The rotation matrix is used to rotate about the origin of the principle coordinate axes, in the decision space.

Algorithm 1 Algorithm for generating a uniformly random rotation matrix. R is an $m \times m$ rotation matrix, where R_i is the ith row of R, and $R^{(j)}$ is the jth column of R. $N(0,1)$ is a normal distribution with a mean of 0 and variance of 1.

for $i = 0$ to m **do**
  for $j = 0$ to m **do**
    $R_i^{(j)} = N(0,1)$
  end for
  for $j = 0$ to m **do**
    $R_i^{(j)} = \frac{R_i^{(j)}}{\|R_i\|}$
  end for
  $\vec{d} = R_i$
  for all j such that $0 \leq j \leq i$ **do**
    $n = \left\|\vec{d}\right\|$
    $p = \vec{d} \cdot R_j$
    for $k = 0$ to m **do**
      $d_k = \frac{d_k - p \cdot R_j^{(k)}}{n^2}$
    end for
  end for
  $R_i = \frac{\vec{d}}{\|\vec{d}\|}$
end for

3.2 Rotated Test Problems

In this section, four rotated test problems are introduced. These problems can be arbitrarily rotated in the decision space. Each of these problems has at least one objective which is non-linear, and through a rotation of the coordinate system, parameter interactions can be introduced.

Problem R1 was first proposed by Deb [1]. It is characterised by a valley in objective f_2. The Pareto-optimal set is situated along the length of this valley as well, and when the problem is subject to a rotation the valley can trap a non-rotationally invariant search from progressing along it, as was explained in Section 2.1. The function f_1 is linear and f_2 is non-linear.

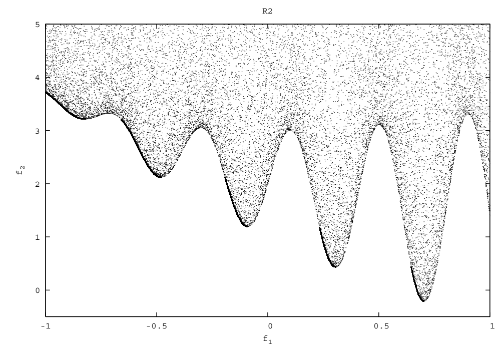

Figure 6: R2 Pareto-optimal front and feasible region

$$\left.\begin{array}{l} f_1(\mathbf{y}) = y_1 \\ f_2(\mathbf{y}) = g(\mathbf{y})h(f_1(\mathbf{y}), g(\mathbf{y})) \\ h(f_1(\mathbf{y}), g(\mathbf{y})) = \exp\left(\frac{-f_1(\mathbf{y})}{g(\mathbf{y})}\right) \\ g(\mathbf{y}) = 1 + 10(m-1) + \sum_{i=2}^{m}\left[y_i^2 - 10\cos(4\pi y_i)\right] \\ \mathbf{y} = \mathbf{Rx},\ -0.3 \leq x_i \leq 0.3,\ for\ i = 1, 2, ..., m \\ |f_1| \leq 0.3 \end{array}\right\} \quad \text{R1}$$

Problem R2 is similar to the ZDT3 problem, and has a Pareto-optimal front which is not continuous. R2 presents a difficulty to an optimization algorithm, because it has to locate a number of discontinuous Pareto-optimal fronts, and maintain solutions in each of those fronts. When R2 is rotated, an optimization algorithm which only searches independently along the principle coordinate axes will generate non-dominated solutions which skew away significantly from the Pareto-optimal front. The reason for this behaviour is that perturbed solutions have to travel quite far along the principle coordinate axes before an independent perturbation can generate a solution which dominates the current non-dominated set. The function f_1 is linear and f_2 is non-linear. The g function is not multi-modal over the specified range. The feasible space and Pareto-optimal front of this problem is shown in Figure 6.

$$\left.\begin{array}{l} f_1(\mathbf{y}) = y_1 \\ f_2(\mathbf{y}) = g(\mathbf{y})h(f_1(\mathbf{y}), g(\mathbf{y})) \\ h(f_1(\mathbf{y}), g(\mathbf{y})) = 1.0 + \exp\left(\frac{-f_1(\mathbf{y})}{g(\mathbf{y})}\right) + \\ \left(\frac{f_1(\mathbf{y}) + 1.0}{g(\mathbf{y})}\right)(\sin(5\pi f_1(\mathbf{y}))) \\ g(\mathbf{y}) = 1 + 10(m-1) + \sum_{i=2}^{m}\left[y_i^2 - 10\cos(\pi y_i)\right] \\ \mathbf{y} = \mathbf{Rx},\ -1.0 \leq x_i \leq 1.0,\ for\ i = 1, 2, ..., m \\ |f_1| \leq 1.0 \end{array}\right\} \quad \text{R2}$$

Decision space variables which increment at a regular interval, evaluate with non-regular intervals in the objective space on Problem R3, making it hard to find a uniform distribution along the Pareto-optimal front. The density of

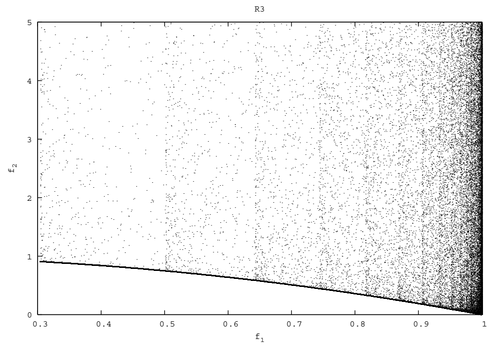

Figure 7: R3 Pareto-optimal front and feasible region

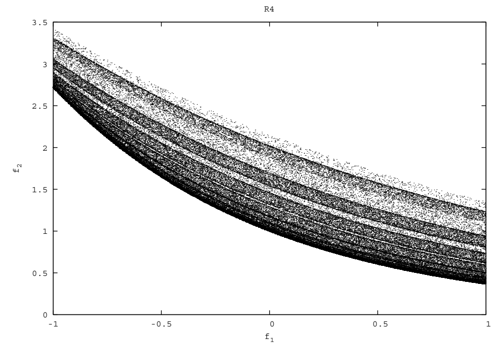

Figure 8: R4 Pareto-optimal front and feasible region

solutions is lower towards lower f_1 values. Problem R3 is similar to the ZDT6 problem. f_1 and f_2 are non-linear functions. The g function is not multi-modal over the specified range. The feasible space and Pareto-optimal front of this problem is shown in Figure 7.

$$\left.\begin{array}{l} f_1(\mathbf{y}) = 1.0 - \exp(2.0y_1)\sin^6(6\pi y_1)/9.0 \\ f_2(\mathbf{y}) = g(\mathbf{y})h(f_1(\mathbf{y}), g(\mathbf{y})) \\ h(f_1(\mathbf{y}), g(\mathbf{y})) = 1.0 - \left(\dfrac{f_1(\mathbf{y})}{g(\mathbf{y})}\right)^2 \\ g(\mathbf{y}) = 1 + 10(m-1) + \displaystyle\sum_{i=2}^{m}\left[y_i^2 - 10\cos(\pi y_i)\right] \\ \mathbf{y} = \mathbf{R}\mathbf{x},\ -1.0 \le x_i \le 1.0,\ for\ i = 1,2,...,m \\ 0.3 \le f_1 \le 1.0 \end{array}\right\} \quad \text{R3}$$

Problem R4 is based on the Schwefel function [12], where a local front is located far from the global minimum. Points are easily trapped by this deceptive front. R4 is difficult, and objective f_2 is characterised by a number of valleys, including the highly deceptive valleys far from the true Pareto-optimal front. These valleys correspond to the modalities generated by function g. Each of these valleys can trap points in a sub-optimal non-dominated front. The feasible space and Pareto-optimal front of this problem is shown in Figure 8.

$$\left.\begin{array}{l} f_1(\mathbf{y}) = y_1 \\ f_2(\mathbf{y}) = g(\mathbf{y})h(f_1(\mathbf{y}), g(\mathbf{y})) \\ h(f_1(\mathbf{y}), g(\mathbf{y})) = \exp\left(\dfrac{-f_1(\mathbf{y})}{g(\mathbf{y})}\right) \\ g(\mathbf{y}) = 1.0 + 0.015578(m-1.0) + \\ \displaystyle\sum_{i=2}^{m}(y_i^2 - 0.25(y_i\sin(32.0\sqrt{|y_i|}))) \\ \mathbf{y} = \mathbf{R}\mathbf{x},\ -1.0 \le x_i \le 1.0,\ for\ i = 1,2,...,m \\ |f_1| \le 1.0 \end{array}\right\} \quad \text{R4}$$

4. EXPERIMENTS

An algorithm for optimizing a multi-objective problem which is not aligned with the principle coordinate axes must have the property of rotational invariance. Secondly, like other EMO algorithms it must maintain good coverage and spread as well as convergence towards the Pareto-optimal front.

We have used a rotationally invariant Non-dominated Sorting Differential Evolution algorithm with Directional Convergence and Spread (NSDE-DCS) for the purposes of contrasting the performance of NSGA-II on the proposed rotated problems [5]. Each of the test problems described in the previous section used a 10 dimensional decision space for this study. For the NSDE-DCS, F was set to 0.8 and K was set to 0.4. These settings were also employed in a preliminary study of rotational invariance [4]. The NSGA-II[1] experiments used a mutation rate of 0.1 and crossover rate of 0.9. η_c and η_m are parameters within the NSGA-II which control the distribution of the crossover and mutation probabilities and were assigned values of 10 and 50 respectively. For both algorithms, a population size of 100 individuals was employed, and 50 runs of each algorithm were conducted, for each test problem.

Experiments were conducted on each of the test problems. Rotations were performed in the decision space, using a random uniform rotation matrix generated using the technique described in Section 3.1. In 10-dimensions there are 45 planes of rotation, introducing parameter interactions between each parameter with every other. A new random uniform rotation matrix was generated for each run of each algorithm.

5. DISCUSSION

The results of our experiments are presented in Figures 9 and 10. For each run of the NSGA-II and NSDE-DCS we have plotted the final non-dominated solutions set after 1000 generations. The purpose of presenting these plots is to demonstrate the type of behaviour one can expect from two types of algorithms; a rotationally invariant EMO and a non-rotationally invariant EMO, on a variety of rotated multi-objective problems. The plots demonstrate the difficulty in convergence towards the Pareto-optimal front, as well as the discovery of non-dominated solutions which are not Pareto-optimal, which occurs when the EMO algorithm is not rotationally invariant.

From Figure 9 it is apparent that the majority of runs, of the rotationally invariant NSDE-DCS, converge and cover

[1]The variant of NSGA-II used in this study is the original NSGA-II available from the Kanpur Genetic Algorithms Laboratory site at `http://www.iitk.ac.in/kangal/`

the Pareto-optimal front of R1. For the discontinuous R2, and the non-uniformly mapped R3, some runs closely approximated the Pareto-optimal front more or less. The difference in performance, with the NSGA-II, is striking on each of these problems. In Figure 10 it is apparent that of the 50 runs of the NSGA-II algorithm, not a single run succeeded in covering the Pareto-optimal front of R1. For R2, and R3, only a few runs managed to find near Pareto-optimal solutions. Non-dominated solutions had a tendency to skew away from the Pareto-optimal front on R2 when the NSGA-II was employed, which is consistent with the description in Section 3. In particular, complete coverage of the Pareto-optimal front of R2 was not possible with NSGA-II over the 50 runs. Problem R3 was the only proposed problem which had a non-linear f_1 and f_2 function, and it presented significant difficulties for NSGA-II with respect to convergence to the Pareto-optimal front, and the spread of solutions generated.

R4 is a highly deceptive multimodal problem, with a local front which is close to the Pareto-optimal front with respect to fitness. This local front maps to a region of the decision space which is far from the global optimal Pareto-optimal front. Over the 50 runs of each algorithm, the NSDE-DCS demonstrated far better coverage of this front than the NSGA-II, although both demonstrated difficulty in converging to the Pareto-optimal front, and had a tendency to become trapped in a local front.

6. CONCLUSION

EMO algorithms should be invariant under a coordinate rotation in order to efficiently optimize complex multi-objective problems with many parameter interactions. Many current results reported in the literature are with respect to problems which do not exhibit complex parameter interactions, although such interactions are characteristic of many real-world problems [7].

We have demonstrated four problems with different characteristics which maintain the same fitness landscape under an arbitrary rotation in the decision space. We have compared NSGA-II with a rotationally invariant algorithm, NSDE-DCS, in order to contrast the difference in behaviour. It is apparent that a rotationally invariant scheme, such as the NSDE-DCS, demonstrates superior performance on problems with significant parameter interactions, compared with the NSGA-II. The primary reason for the poor performance of NSGA-II on the problems presented in this paper, is that NSGA-II uses a non-rotationally invariant crossover operator.

Although we have demonstrated a small range of problems, many others are conceivable, including problems with more than two objectives. Increasing the number of objectives results in rotated problems which are potentially more challenging for non-rotationally invariant EMO algorithms, because there will be more non-dominated solutions to choose from which are not optimal. It is relatively easy to find such non-dominated solutions, but it is harder to find the Pareto-optimal set.

This paper has proposed a starting point for re-evaluating many current EMO algorithms, and new algorithms which have yet to be proposed. with respect to the important topic of rotational invariance.

7. REFERENCES

[1] K. Deb. Multi-objective genetic algorithms: Problem difficulties and construction of test problems. *Evolutionary Computation*, 7(3):205–230, 1999.

[2] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multi-objective genetic algorithm: Nsga-ii. *IEEE Trans. Evol. Comput.*, 6(2):182–197, 2002.

[3] S. Huband, L. Barone, L. While, and P. Hingston. A scalable multi-objective test problem toolkit. In *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005, Lecture Notes in Computer Science*, volume 3410, pages 280–295, Mexico, 2005.

[4] A. Iorio and X. Li. Solving rotated multi-objective optimization problems using differential evolution. In *AI 2004: Advances in Artificial Intelligence: 17th Australian Joint Conference on Artificial Intelligence, Cairns, Australia*, page 861, Heidelberg, 2004.

[5] A. Iorio and X. Li. Incorporating directional information within a differential evolution algorithm for multi-objective optimization. In *Proceedings of the 2006 Genetic and Evolutionary Computation Conference (GECCO-06)*, 2006.

[6] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable test problems for evolutionary multi-objective optimization. Technical Report TIK-Report No. 112, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH), Zurich, 2001.

[7] J. Knowles. Parego: A hybrid algorithm with on-line landscape approximation for expensive multi-objective optimization problems. *IEEE Trans. Evol. Comput.*, 10(1):50–66, 2005.

[8] J. Knowles and D. Corne. Instance generators and test suites for the multi-objective quadratic assignment problem. In *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003, Lecture Notes in Computer Science*, volume 2632, pages 295–310, Portugal, 2003.

[9] T. Okabe and Y. Jin. On test functions for evolutionary multi-objective optimization. In *Parallel Problem Solving from Nature - PPSN VIII, Lecture Notes in Computer Science*, volume 3242, pages 792–802, Birmingham, 2004.

[10] K. Price. *Differential Evolution*, chapter 2, pages 79–108. McGraw-Hill, London UK, 1999.

[11] T. Robic. Performance of demo on new test problems: A comparison study. In *Proceedings of the Fourteenth International Eletrotechnical and Computer Science Conference ERK 2005*, pages 121–124, 2005.

[12] R. Salomon. Re-evaluating genetic algorithm performance under coordinate rotation of benchmark functions: A survey of some theoretical and practical aspects of genetic algorithms. *Bio. Systems*, 39(3):263–278, 1996.

[13] E. Zitzler, K. K. Deb, and L. Thiele. Comparison of multi-objective Evolutionary Algorithms: Empirical results. *Evolutionary Computation*, 8(2):173–195, 2000.

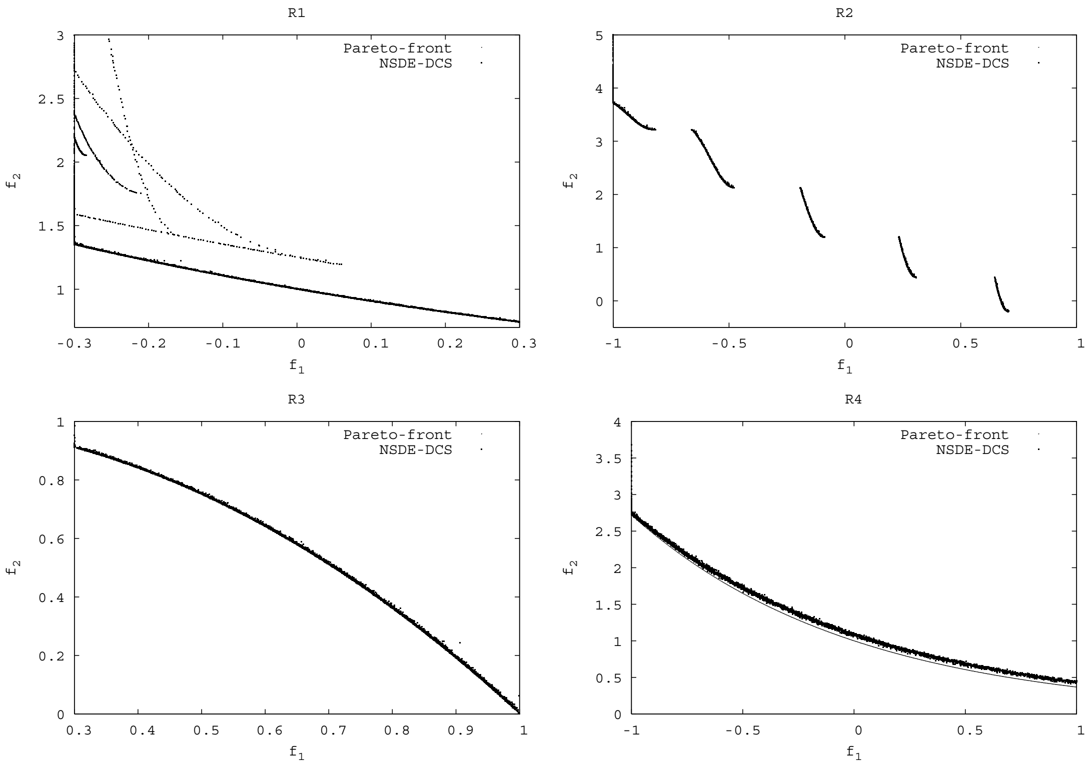

Figure 9: 50 runs of NSDE-DCS (Generation 1000)

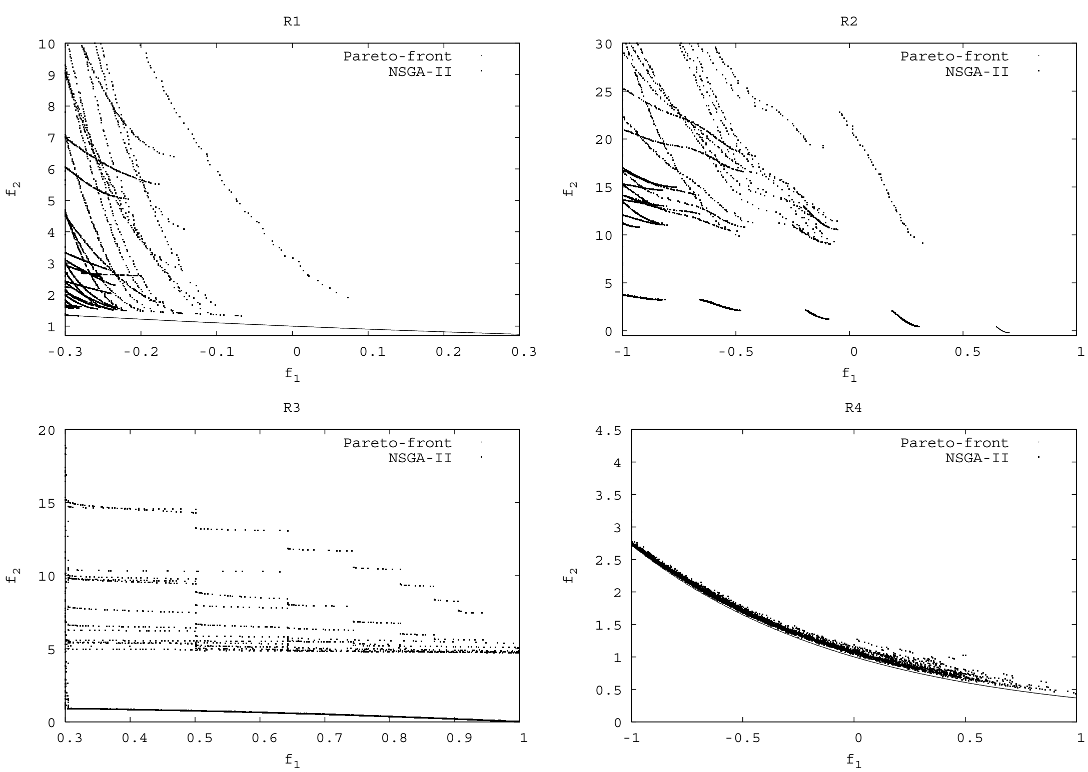

Figure 10: 50 runs of NSGA-II (Generation 1000)