

**A novel preference articulation operator for the
Evolutionary Multi-Objective Optimisation of classifiers in
concealed weapons detection**

ROSTAMI, Shahin, O'REILLY, Dean, SHENFIELD, Alex <<http://orcid.org/0000-0002-2931-8077>> and BOWRING, Nicholas

Available from Sheffield Hallam University Research Archive (SHURA) at:

<https://shura.shu.ac.uk/8771/>

This document is the Accepted Version [AM]

Citation:

ROSTAMI, Shahin, O'REILLY, Dean, SHENFIELD, Alex and BOWRING, Nicholas (2015). A novel preference articulation operator for the Evolutionary Multi-Objective Optimisation of classifiers in concealed weapons detection. *Information Sciences*, 295, 494-520. [Article]

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

A Novel Preference Articulation Operator for the Evolutionary Multi-Objective Optimisation of Classifiers in Concealed Weapons Detection

Shahin Rostami^a, Dean O'Reilly^a, Alex Shenfield^b, Nick Bowring^a

^a*School of Engineering, Manchester Metropolitan University, Manchester, M1 5GD, United Kingdom*

^b*Department of Engineering and Mathematics, Sheffield Hallam University, Sheffield, S1 1BW, United Kingdom*

Abstract

The incorporation of decision maker preferences is often neglected in the Evolutionary Multi-Objective Optimisation (EMO) literature. The majority of the research in the field and the development of EMO algorithms is primarily focussed on converging to a Pareto optimal approximation close to or along the true Pareto front of synthetic test problems. However, when EMO is applied to real-world optimisation problems there is often a decision maker who is only interested in a portion of the Pareto front (the Region of Interest) which is defined by their expressed preferences for the problem objectives. In this paper a novel preference articulation operator for EMO algorithms is introduced (named the Weighted Z-score Preference Articulation Operator) with the flexibility of being incorporated *a priori*, *a posteriori* or *progressively*, and as either a primary or auxiliary fitness operator. The Weighted Z-score Preference Articulation Operator is incorporated into an implementation of the Multi-Objective Evolutionary Algorithm Based on Decomposition (named WZ-MOEA/D) and benchmarked against MOEA/D-DRA on a number of bi-objective and five-objective test problems with test cases containing preference information. After promising results are obtained when comparing WZ-MOEA/D to MOEA/D-DRA in the presence of decision maker preferences, WZ-MOEA/D is successfully applied to a real-world optimisation problem to optimise a classifier for Concealed Weapon Detection, producing better results than previously published classifier implementations.

Keywords: Evolutionary Computation, Evolutionary Multi-Objective Optimisation, Evolutionary Algorithms, Preference Articulation, Concealed Weapon Detection, Metaheuristics

1. Introduction

The optimisation of the accuracy and efficiency of classifiers in pattern recognition is a complex problem that is often poorly understood. For example, whilst numerous techniques

Email address: s.rostami@mmu.ac.uk (Shahin Rostami)

exist for the optimisation of weights in Artificial Neural Networks (ANNs) (such as the Widroff-Hoff least mean squares algorithm and back propagation), there do not exist any hard and fast rules for choosing the structure of an ANN - in particular for choosing both the size (in term of number of neurons) and number of hidden layers used in the network. However, this internal structure is one of the key factors in determining the efficiency of the network and the accuracy of the classification. In recent years there has been some interest in using soft computing techniques such as Evolutionary Algorithms (EAs) to provide a solution to this problem [22], focussing on evolving the structure of an ANN to solve function approximation problems. However, complex classification problems often involve trade-offs between classification objectives that are not well suited to this kind of single objective approach.

One approach to solving complex engineering problems is to use Evolutionary Multi-Objective Optimisation (EMO) algorithms to address each of the conflicting objectives simultaneously. Typically, these EMO algorithms are run non-interactively with a decision maker setting the initial parameters of the algorithm and then analysing the results and the end of the execution process (which can often take hours or days to complete). This approach has been common since the late 1990s [2, 34, 36] and will lead to a set of potential solutions distributed across the whole trade-off surface. Whilst this is often appropriate for problems with a small number of objectives, in real world problems that involve the consideration of many objectives this trade-off surface can be very large. In these cases, the decision maker is usually more interested in a sub-region of this solution space that satisfies some domain specific criteria. However, this can be complicated further by a lack of *a priori* knowledge about what trade-offs are achievable. To overcome these problems, progressive preference articulation methods have been proposed that take into account decision maker preferences (such as [4]) but these are frequently difficult to integrate with current state-of-the-art EMO algorithms, and the incorporation of user preferences is frequently disregarded in the EMO literature [6].

In this research the authors present the optimisation of the ANN architecture for a two objective problem and a five objective problem. The two objective optimisation is performed on an ANN that is classifying the radar signals into two groups which are *threat* and *non-threat*. The five objective optimisation is ambitious in the sense that it attempts to optimise the architecture of a number of ANNs each of which are trained to detect a specific threat item. This multi-objective optimisation is the more difficult of the two problems but will give a greater level of information to the user of the detection system. This will allow the security forces to react to specific threats in a more controlled manner as they will know the type of threat presented.

This paper introduces a novel method of progressive preference articulation in EMO algorithms which can provide improved performance in both the execution speed of the algorithm and in the quality of the solutions the algorithm produces. This method is then integrated into a state-of-the-art EMO algorithm and applied both to current benchmark optimisation problems from the literature and to a real-world classification problem in the field of concealed weapon detection. Section 1.1 introduces EAs and discusses their suitability to multi-objective optimisation. Section 2 describes the implementation of the progressive

preference articulation method and discusses its effectiveness when applied to some deceptive test functions from the literature. Section 3 provides a full statistical analysis of the results of the integration of the proposed novel progressive preference articulation operator with a state-of-the-art EMO algorithm for two suites of benchmark test functions from the literature. Section 4 shows the effectiveness of the proposed framework in the optimisation of ANN structures when applied to a complex real-world classification problem in concealed weapon detection, the results of which are compared against existing published work. The main results from Section 3 and Section 4 are then summarised in Section 5 and some conclusions about the use of and the effectiveness of the proposed framework and its suitability to the optimisation of ANN structures for classification are drawn.

1.1. Evolutionary algorithms

EAs are a powerful class of stochastic optimisation techniques that incorporate some of the principles of natural selection and population genetics to converge towards global optima [19]. They provide an iterative and population-based approach to optimisation that is capable of both exploring the search space of a problem and exploiting promising solutions found in previous generations. Typically the exploration of the search space is performed by using variation operators (such as mutation) that introduce an element of stochasticity into the optimisation process and aim to prevent premature convergence to local optima. In contrast, exploitation of promising solutions from previous generations is performed using a selection operator (and in part, recombination operators) that ensures preference is given to solutions that are considered fittest from the previous generation.

The robustness of EAs to multi-modal search landscapes containing many local optima (any other difficulties present in multi-objective search spaces) and the direct use of objective function information (rather than auxiliary knowledge such as derivative information) ensures that EAs are effective when applied to many problem types in which conventional optimisation methods may have difficulty. In addition, their population-based nature helps ensure that EAs are resilient when faced with noisy search spaces, as each generation contains more information about the shape of the fitness landscape than would be available to conventional, non-population based methods such as hill-climbing [26].

1.2. Multi-objective optimisation and EAs

The general form of a multi-objective optimisation problem can be described by an objective vector f and a corresponding set of design variables v ; as can be seen in Equation 1. Note that here minimisation can be assumed with no loss of generality.

$$\min_f(v) = (f_1(v), f_2(v), f_3(v), \dots, f_n(v)) \quad (1)$$

Multi-objective optimisation problems often involve conflicts between multiple objectives, and as a result it is unlikely that there exists a single optimal solution. Instead, the solution of a multi-objective optimisation problem often consists of a candidate population of Pareto optimal points - where any improvement in one objective will result in the degradation of one or more of the other objectives.

A set of non-dominated solutions¹ generated by the optimiser is known as an *approximation set* [43] and can be characterised in three key areas [29]. These are illustrated graphically in Figure 1, and listed below.

- Proximity. This tells the decision maker how close the approximation set is to the true Pareto front. An ideal approximation set should be as close as possible to the true Pareto front.
- Diversity. This characterises the distribution of the approximation set both in the extent and uniformity of that distribution. The ideal approximation set should be uniformly distributed across the trade-off surface of the problem.
- Pertinence. This criteria measures the relevance of the approximation set to the decision maker. Ideally the approximation set should contain a number of solutions that potentially satisfy the decision maker’s expressed preferences.

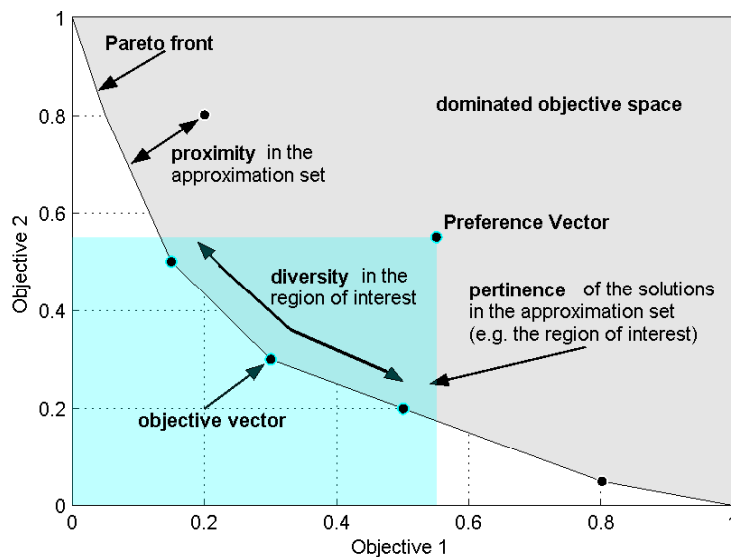


Figure 1: Proximity, diversity, and pertinence characteristics in an approximation set in bi-objective space.

Conventional multi-objective optimisation techniques often fail to satisfy these criteria. For example, the goal-attainment method [18] and the weighted-sum method [24] both only provide single solutions to the optimisation problem - thus failing to provide a diverse distribution of solutions. However, EAs are well suited to this kind of multi-objective optimisation

¹A solution is termed non-dominated if there is no other solutions in the solution set that is superior in all objectives.

since they search a population of candidate solutions and are thus capable of presenting a diverse approximation set to a decision maker [8].

Many theoretical EMO studies only consider a small number of objectives, with most of the published literature focussing on the bi-objective case. However complex real-world engineering problems often require the consideration of a larger number of objectives. This has led to much interest amongst the research community in *many-objective optimisation*². The increased scale of a many-objective optimisation problem means that the *pertinence* of the approximation set presented to a decision maker is especially important. The global trade-off surface for a problem with many conflicting objectives frequently contains many Pareto optimal solutions, the majority of which may not be in the decision maker's Region of Interest (ROI) [29].

1.3. Preference articulation and decision making

The role of the decision maker in the EMO process is usually to choose a single compromise solution from the approximation set presented to them. Although there may be a potentially infinite number of Pareto optimal solutions in the global trade-off surface, in practice the decision maker will usually only be interested in a small subset of these. Therefore, allowing the decision maker to focus the optimisation process on relevant areas of the search space both increases the efficiency of the search effort and reduces the amount of irrelevant information the decision maker has to consider [15].

The preferences of a decision maker can be incorporated into the optimisation process in three ways:

- *A posteriori*
- *A priori*
- Progressively

A posteriori methods of preference articulation involve the decision maker selecting a compromise solution from the global approximation set of Pareto optimal solutions found at the end of the optimisation process, whilst *a priori* and progressive preference articulation methods aim to achieve a good representation of the trade-off surface in the ROI of the decision maker. The key advantage of *a priori* and progressive preference articulation methods is the reduction in the size of the search space explored by the optimiser because the search is focussed on a sub-set of the global trade-off surface.

In *a priori* articulation of preferences the decision maker expresses their preferences before the start of the optimisation process. However, often the decision maker may not be sure of their preferences prior to optimisation and, by stating their preferences *a priori*, the decision maker may not investigate some areas of the search space that deserve attention. A better method is often progressive articulation of preferences, which enables the decision

²The phrase *many-objective* has been used in the operations research community to refer to problems with more than the standard two or three objectives [14].

maker to alter their preferences during the search and thus incorporate information that only becomes available during the search process (such as the exact nature of trade-offs between objectives).

One of the first schemes for progressive preference articulation in EMO algorithms was introduced by [17], and extended the Pareto-based ranking scheme used in the Multiple Objective Genetic Algorithm (MOGA) [16] to allow preferences to be expressed throughout the run of a EMO algorithm. These preferences were then used in a modified version of dominance which combines the concept of Pareto optimality with a preference operator to rank the candidate solutions according to both preference information and Pareto-dominance. This progressive preference articulation method has been used in a wide variety of engineering applications such as the optimisation of robust control strategies for gasifier power plants [20] and the design of lateral stability controllers for aircraft [31].

2. Weighted Z-score preference articulation

Weighted Z-score (WZ) preference articulation is a novel method of preference articulation based around the use of z-scores³ (or standard scores) from statistics. Traditional z-score calculations are performed by subtracting the population mean from a datum and then dividing the result by the population standard deviation as can be seen in Equation 2. Calculating the z-score in statistics requires knowing the population parameters and not just the parameters of a sample, which is often seen as unrealistic in typical statistics; however this is not an issue in EMO as it is possible to have a complete representation of the population at each generation.

$$z = \frac{(x - \mu)}{\sigma} \quad (2)$$

For the z-score to be useful for preference articulation, some modifications are made to the way z is calculated. Instead of using the population mean and population standard deviation to calculate z , the preference information⁴ that has been expressed by the decision maker is used (as can be seen in Equation 3) where ρ_m is the goal for a corresponding objective value x_{mn} , and N is the number of solutions in the population.

$$z_{mn} = \frac{(x_{mn} - \rho_m)}{\sqrt{\frac{\sum_{n=1}^N (x_{mn} - \rho_m)^2}{N}}} \quad (3)$$

This will enable the calculation of z_{mn} for the objective values of each candidate solution in an approximation set, resolving the number of standard deviations each solution is from the decision makers expressed ROI, which will be a positive value when it is outside the ROI, and negative when within the ROI. Once z_{mn} is calculated for every objective value of a solution, the z_{mn} values are aggregated into a single fitness value using Equation 4.

³The number of standard deviations a datum is above or below the mean of its data-set.

⁴A ROI can be defined by a preference vector containing goals for each objective, where if all objective values of a solution satisfy the corresponding objective goal it is considered within the ROI.

$$V_n = \frac{\sum_{m=1}^M z_{mn}}{M} \quad (4)$$

As a demonstration, the CEC09 competition winning MOEA/D-DRA [40] EMO is executed for five generations to generate an initial population for the ZDT1 synthetic test problem from the ZDT test suite. Using Equation 4 to calculate V_n for each candidate solution, it is possible to order the initial population by the aggregated number of standard deviations from the ROI (0.2 for objective 1 and 1 for objective 2 for this example) and then select a number of solutions (five for this example) to exploit for the next generation. An illustration of this example can be seen in Figure 2.

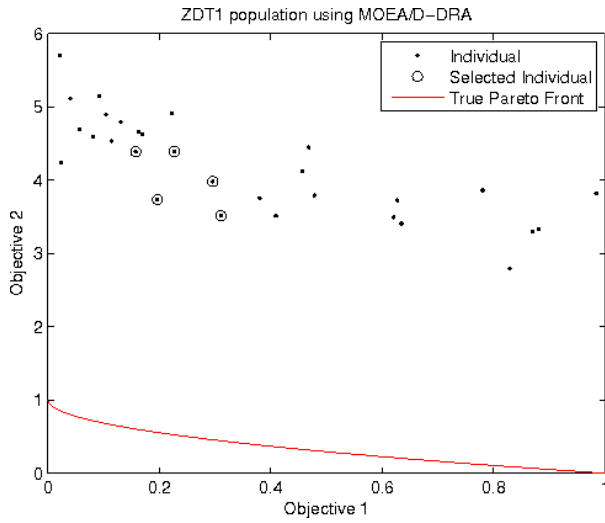


Figure 2: Basic Z-score preference articulation applied to an initial population generated by MOEA/D-DRA for the synthetic test problem ZDT1.

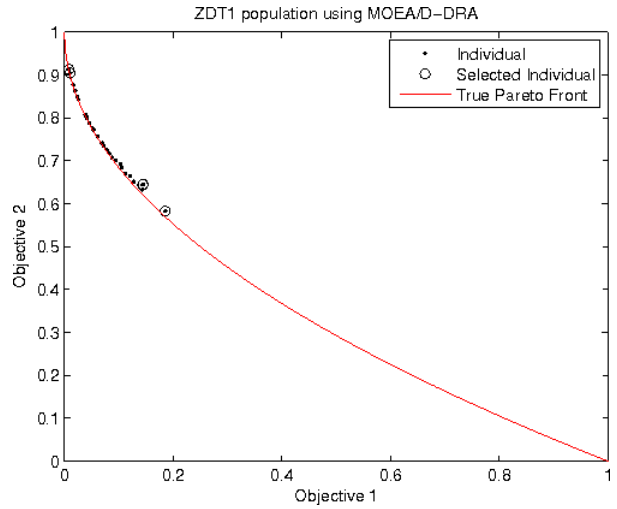


Figure 3: Population generated by MOEA/D-DRA combined with basic Z-score preference articulation after 2000 function evaluations for the synthetic test problem ZDT1.

The Z-scores calculated using the simple method in Equation 3 can be used in an EMO as either a replacement or addition to the fitness scheme used for selection, to focus the search towards and then within the ROI expressed by the decision maker. This works well for test problems where there is a low number of objectives and the Pareto set is not complicated in shape as seen in Figure 3, however its effectiveness is reduced when this basic method is applied to a more complicated problem with a higher number of objectives. To demonstrate this, the same basic Z-score preference articulation method is applied to an initial population generated by executing MOEA/D-DRA for five generations on a five-objective instance of the WFG5 synthetic test problem from the WFG tool-kit, this is illustrated in Figure 4.

This demonstration illustrates that the basic method of Z-score preference articulation falls into the trap of selecting the solutions with the overall lowest $z_m n$ values. This results in the minimisation of objectives 1 to 4 as they are already below or close in proximity to the corresponding preference value ρ_m , yielding negative or low positive z values, and

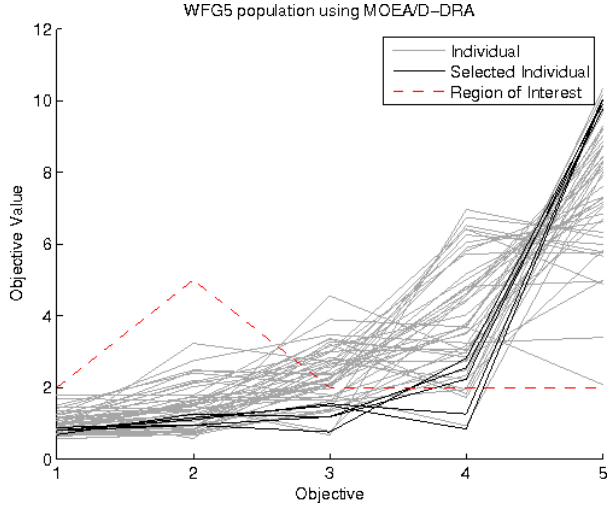


Figure 4: Basic Z-score preference articulation applied to an initial population generated by MOEA/D-DRA for the synthetic test problem WFG5.

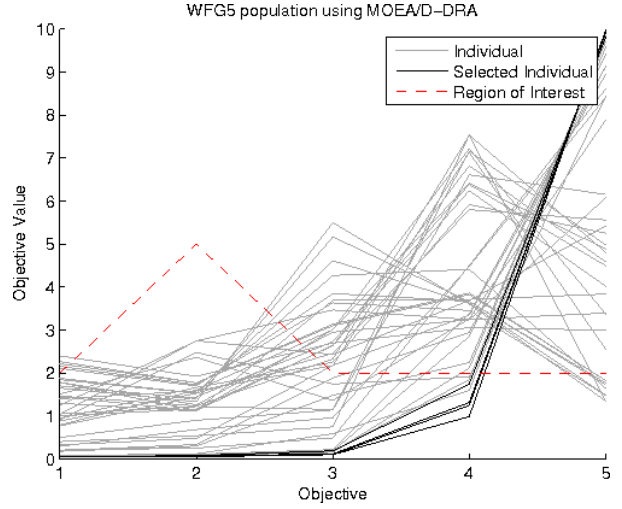


Figure 5: Population generated by MOEA/D-DRA combined with basic Z-score preference articulation after 2000 function evaluations for the synthetic test problem WFG5.

due to the conflicting objectives in WFG5, the associated 5th objective value for these selected solutions are the furthest away from the specified ROI. This issue is amplified from generation to generation and the result after 2000 function evaluations with the ROI specified as $\{2, 5, 2, 2, 2\}$ has been illustrated in Figure 5, where it can be observed that although objectives 1-4 have been further optimised, there are still no individuals within the desired ROI due to objective 5.

To solve this ineffectiveness at higher numbers of objectives and on more difficult problems, a two-phase preference articulation operator has been developed. The first phase (W-phase) focusses the search on bringing all objectives closer in proximity to the desired ROI using an absolute and weighted z values. The second phase (Z-phase) takes effect once a criteria for the number of solutions required within the ROI has been met, this phase uses the basic z values demonstrated earlier for minimisation within the ROI.

The mathematical procedure for the WZ preference articulation operator in its entirety is described herein. M defines the number of problem objectives whilst N defines the population size. X is an M by N matrix of entries x_{mn} , where every x_{mn} refers to a solution's objective value:

$$X_n = \langle x_{1n}, x_{2n}, \dots, x_{Mn} \rangle \quad (5)$$

Z is an M by N matrix of entries z_{mn} , where every z_{mn} refers to the result of the z-score preference articulation operator applied to a corresponding objective value x_{mn} :

$$Z_n = \langle z_{1n}, z_{2n}, \dots, z_{Mn} \rangle \quad (6)$$

To calculate Z , a preference vector P of M entries must be defined, where every entry ρ_m refers to the goal which the corresponding objective values x_m must satisfy:

$$P = \langle \rho_1, \rho_2, \dots, \rho_M \rangle \quad (7)$$

S is an M by N matrix of entries s_{mn} where every s_{mn} refers to a logical value indicating whether the corresponding objective value x_{mn} has satisfied the corresponding goal ρ_m ($x_{nm} \leq \rho_m$):

$$S_n = \langle s_{1n}, s_{2n}, \dots, s_{Mn} \rangle \quad (8)$$

where s_{mn} is calculated using:

$$s_{mn} = \begin{cases} 1, & \text{if } x_{mn} \leq \rho_m \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

Φ is a vector of N entries, where every ϕ_n refers to a logical value indicating whether all entries of P have been satisfied by a solution X_n .

$$\Phi = \langle \phi_1, \phi_2, \dots, \phi_N \rangle \quad (10)$$

where ϕ_n is calculated by the product of the entries of S_n :

$$\phi_n = \prod_{m=1}^M s_{mn} \quad (11)$$

The scalar Ψ refers to the number of solutions X_n in the population which have satisfied the preference vector P :

$$\Psi = \sum_{n=1}^N \phi_n \quad (12)$$

T defines the required number of solutions which satisfy the preference vector before the search changes phase. Whilst $\Psi < T$ the W -phase of the WZ preference articulation operator takes effect. In this phase, the weighting $(1 - \frac{1}{M})$ is only applied to the z_{mn} value if m corresponds to the entry of Ω with the lowest value. ω_m refers to the number of solutions in the population that have satisfied the corresponding ρ_m :

$$\Omega = \langle \omega_1, \omega_2, \dots, \omega_M \rangle \quad (13)$$

ω_m is the sum of columns M in the matrix S and is calculated using:

$$\omega_m = \sum_{n=1}^N s_{mn} \quad (14)$$

With the entries of Ω calculated, the M by N matrix of weighted scores E can be defined as:

$$E_n = \langle \epsilon_{1n}, \epsilon_{2n}, \dots, \epsilon_{mN} \rangle \quad (15)$$

where the corresponding weighted score ϵ_{mn} for each objective value x_{mn} can be calculated using:

$$\epsilon_{mn} = \begin{cases} z_{mn} \left(1 - \frac{1}{M}\right) & \text{if } f(\omega_m, S_{mn}) = 0 \\ z_{mn} & \text{otherwise.} \end{cases} \quad (16)$$

where z_{mn} and ω_m are first normalised to real values between 0 and 1:

$$z_{mn} = f(|z_{mn}|, |Z_m|) \quad (17)$$

using the function $f(k, K)$ where:

$$f(k, K) = \frac{k - \min(K)}{\max(K - \min(K))} \quad (18)$$

The initial calculation of z_{mn} is the same in both phases (W-phase and Z-phase) and is defined in Equation 3. The final score W_n of a single solution is the aggregation of the corresponding ϵ_{mn} entries:

$$W_n = \frac{\sum_{m=1}^M \epsilon_{mn}}{M} \quad (19)$$

This two-phase method attempts to move the search towards the production of solutions that are close in proximity to the ROI and within it, but does not attempt to minimise the solutions beyond the edges of the ROI. When the number of solutions within the ROI has satisfied the threshold ($\Psi \geq T$) the Z-phase takes effect. This phase uses Equation 3 to calculate Z_n and then Equation 4 to aggregate the scores into the scalar V_n , this is because there are adequate solutions (defined by T) that have satisfied all entries of P . These solutions can then be further minimised within the ROI. A full example of both phases of the WZ preference articulation operator is available in Section 2.1.

When the WZ preference articulation operator algorithm is applied to the same initial population previously generated for WFG5 by MOEA/D-DRA, solutions that are closest to the ROI are selected with weighted preference for solutions with objectives which have not yet been satisfied. This has been illustrated in Figure 6, where it can be seen that solutions with worse values for objective 1-4 have been selected in order to exploit their useful genetic information to bring objective 5 closer to and within the ROI.

The second phase of the search is activated when the threshold T has been met or exceeded, for this example T is set to 5 solutions (one tenth of the population size) and executed for the same 2000 function evaluations as before, the results of the search have been illustrated in Figure 7.

The results from this experiment show that 25 solutions have been found within the desired ROI, further investigation has shown that threshold T of 5 solutions was satisfied at

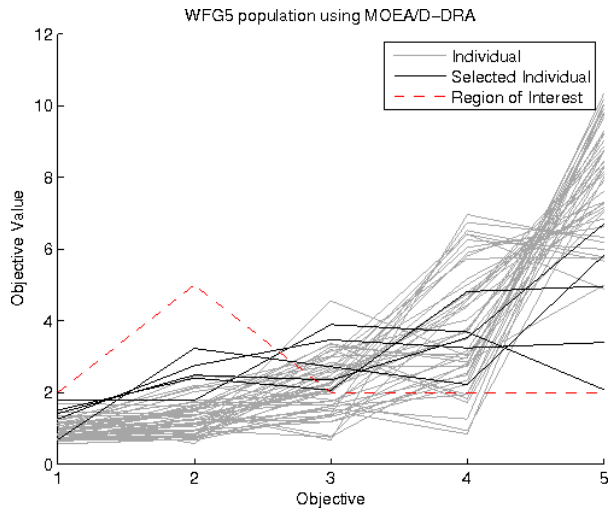


Figure 6: WZ preference articulation applied to an initial population generated by MOEA/D-DRA for the synthetic test problem WFG5.

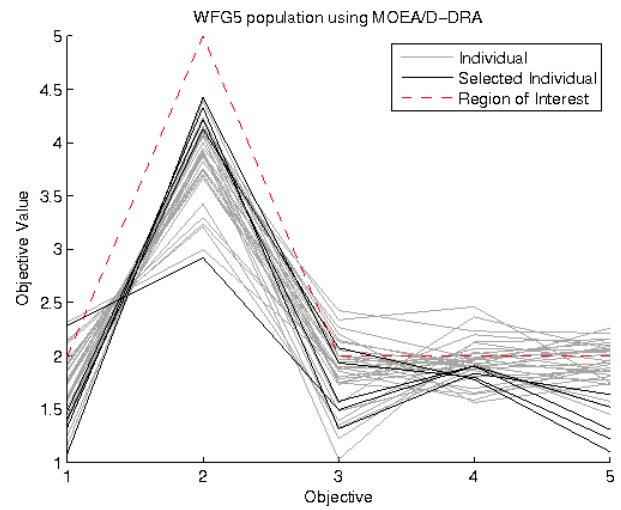


Figure 7: Population generated by MOEA/D-DRA combined with WZ preference articulation after 2000 function evaluations for the synthetic test problem WFG5. 25 solutions have been found in the desired ROI.

just 300 function evaluations which is when the search switched from the W-phase to the Z-phase in the WZ preference articulation operator, the results at 300 function evaluations have been illustrated in Figure 8.

As a preliminary comparison for proof-of-concept MOEA/D-DRA without preference articulation was executed for 2000 function evaluations on WFG5, the results of which are illustrated in Figure 9, for which no solutions were found within the ROI, due to lack of focus toward the desired ROI during the search, this is to be expected in the absence of preference articulation.

The WZ preference articulation operator is algorithm agnostic and can therefore be applied to any EMO algorithm as either primary or secondary sorting criteria for use by a selection operator. As a demonstration, in Figure 10 the WZ preference articulation operator has been combined with an initial population for the ZDT4 synthetic test problem from the ZDT test suite, generated by CMA-PAES [30], an algorithm which uses covariance matrix adaptation for search and adaptive grid archiving for selection and maintenance of a population. After 2000 function evaluations CMA-PAES combined with the WZ preference articulation operator has produced the results illustrated in Figure 11. ZDT4 is a test problem with many deceptive Pareto fronts which makes it computationally expensive for an EMO algorithm to find an approximation set close in proximity to or along the true-front, however when searching toward and within a specified ROI it is possible to cut the computational cost of the search by reducing the search space that is explored, whilst still resolving solution individuals which are of interest to the decision maker.

The WZ preference articulation operator shows promise in finding solutions within a desired ROI by focussing the search and preventing exploration of areas of the search space

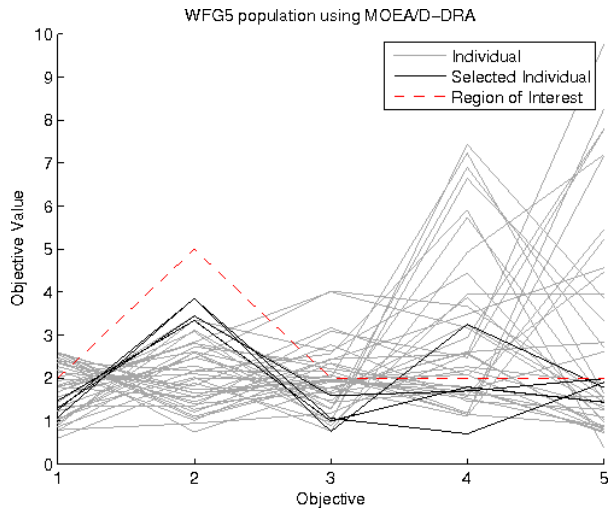


Figure 8: Population generated by MOEA/D-DRA combined with the WZ preference articulation operator after 300 function evaluations for the synthetic test problem WFG5. 6 solutions have been found in the desired ROI.

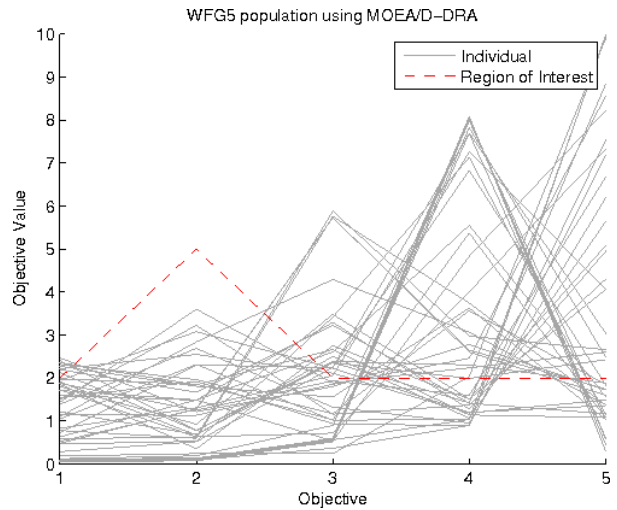


Figure 9: Population generated by MOEA/D-DRA without preference articulation after 2000 function evaluations for the synthetic test problem WFG5. No solutions have been found in the desired ROI.

that will not be of expressed interest to the decision maker, and it can be observed from Figures 7 and 9 that in its absence the performance (in regards to the number of solutions within the ROI) is worse.

The WZ preference articulation operator is a portable and auxiliary operator which can be incorporated into any host⁵ algorithm, therefore its performance in finding solutions within the ROI will ultimately be relative to that of its host and the method of incorporation (whether it is used as the sole measure of fitness or used in conjunction with other fitness operators).

2.1. WZ preference articulation operator worked-example

In this section a complete worked-example of both phases of the WZ preference articulation operator (described in Section 2) is demonstrated. The objective values used in these examples are not the result of any objection function, they have only been selected for demonstration and ease of replication.

2.1.1. The W-Phase

This example assumes a five-objective problem ($M = 5$) with a population of four ($N = 4$) solutions X_n where $T = 2$, this population has been presented in Table 1 and plotted in Figure 12. The preference vector $P = \langle 1, 3, 4, 2, 1 \rangle$ defines the ROI for which the solutions are desired to be within. In order to decide which phase of the WZ preference articulation operator takes effect, Ψ needs to be calculated.

⁵An algorithm that incorporates the WZ preference articulation operator is referred to as the host.

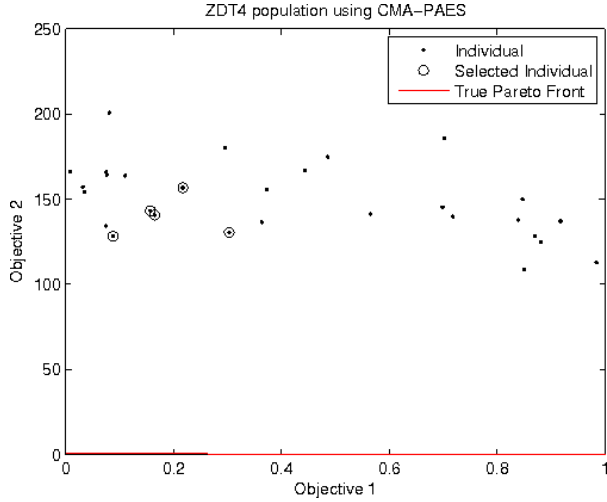


Figure 10: Basic Z-score preference articulation applied to an initial population generated by CMA-PAES for the synthetic test problem ZDT4.

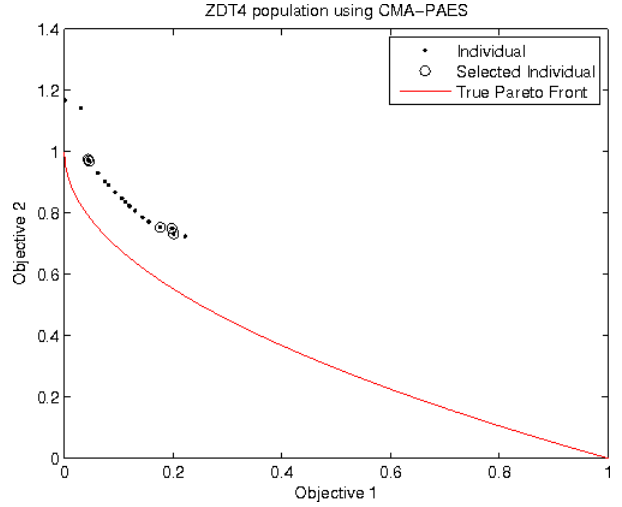


Figure 11: Population generated by CMA-PAES combined with basic Z-score preference articulation after 2000 function evaluations for the synthetic test problem ZDT4.

	x_{1n}	x_{2n}	x_{3n}	x_{4n}	x_{5n}
X_1	0.5	0.5	5.0	2.5	1.5
X_2	0.6	0	5.0	3.0	1.4
X_3	0.5	3.5	4.5	2.5	1.5
X_4	0.8	3.2	4.2	3.0	1.2

Table 1: An example population X of objective values X_{nm} .

	s_{1n}	s_{2n}	s_{3n}	s_{4n}	s_{5n}	$\mathbf{0}$	Ψ
S_1	1	1	0	0	0	0	ϕ_1
S_2	1	1	0	0	0	0	ϕ_2
S_3	1	0	0	0	0	0	ϕ_3
S_4	1	0	0	0	0	0	ϕ_4
Ω	4	2	0	0	0		
$f(\omega, S)$	1	0.5	0	0	0		

Table 2: Logical values of matrix S , vector Ω and it's normalised values, vector Φ and the scalar Ψ calculated for population X .

Table 2 shows that Ψ for the current population against the preference information has been resolved as 0 and because $\Psi < T$ the W-phase of the WZ preference articulation operator takes effect. In order to find out which preference entries ρ_m have been satisfied the least by the population, the entries ω of vector Ω are calculated. The entries of Ω are then normalised to values between 0 and 1 using $f(k, K)$ and it can be seen that ρ_3 , ρ_4 and ρ_5 have the least number of solutions satisfied, therefore the scores for those objectives will receive weighting.

In Table 3 it can be seen that ϵ_{mn} has been calculated for every solution's objective value and has been aggregated per solution into a single score as the scalar W_n . Sorting the entries of W in ascending order resolves (as calculated by the WZ preference articulation operator) the proximity of each solution X_n to P in the objective space. In this example, the solutions

	ϵ_{1n}	ϵ_{2n}	ϵ_{3n}	ϵ_{4n}	ϵ_{5n}	W	
E_1	1.0	0.82	0.8	0	0.8	0.86	W_1
E_2	0.67	1	0.8	0.8	0.53	0.95	W_2
E_3	1	0.11	0.3	0	0.8	0.55	W_3
E_4	0	0	0	0.8	0	0.2	W_4

Table 3: The matrix E containing the weighted scores for the population X , and the aggregated weighted score W for each solution X_n

are ordered (presented in order of ascending proximity to P) $\langle X_4, X_3, X_1, X_2 \rangle$.

In the event that no solutions are found within the ROI, or if $\Psi < T$ is not satisfied throughout the optimisation process, then the WZ preference articulation operator will remain in the W-Phase until the host algorithm meets its termination criteria. The result of the WZ preference articulation operator remaining in the W-Phase is a final approximation set of solutions X_n that are close in proximity to to the preference vector P . This allows the algorithm to still produce an approximation set of feasible solutions that are close in proximity to the decision makers ROI in the scenario where no solutions exist within the ROI.

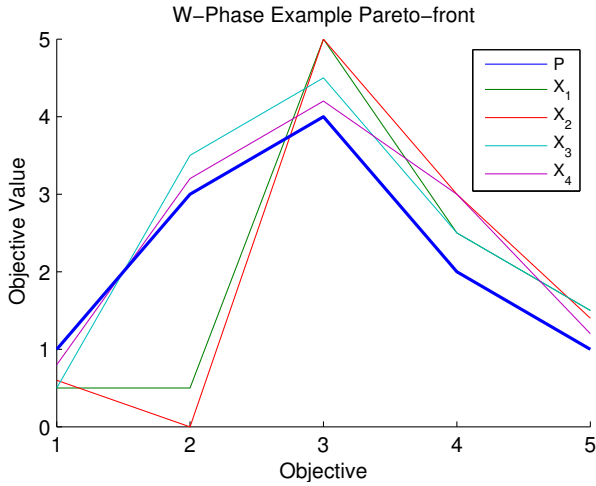


Figure 12: Parallel-coordinate plot of the Population X and preference vector P used in the W-Phase example.

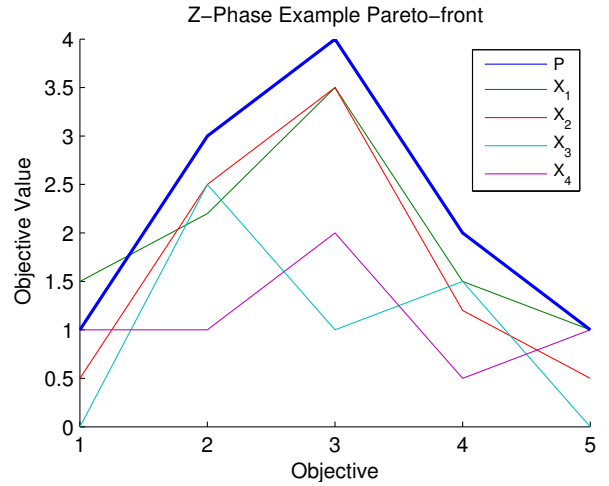


Figure 13: Parallel-coordinate plot of the Population X and preference vector P used in the Z-Phase example.

2.1.2. The Z-Phase

This example also assumes a five-objective problem ($M = 5$) with a population of four ($N = 4$) solutions X_n where $T = 2$, this population has been presented in Table 4 and plotted in Figure 13. The preference vector $P = \langle 1, 3, 4, 2, 1 \rangle$ defines the ROI for which the

solutions are desired to be within. In order to decide which phase of the WZ preference articulation operator takes effect, Ψ needs to be calculated.

	x_{1n}	x_{2n}	x_{3n}	x_{4n}	x_{5n}
X_1	1.5	2.2	3.5	1.5	1.0
X_2	0.5	2.5	3.5	1.2	0.5
X_3	0	2.5	1.0	1.5	0
X_4	1.0	1.0	2.0	0.5	1.0

Table 4: An example population X of objective values X_{nm} .

	s_{1n}	s_{2n}	s_{3n}	s_{4n}	s_{5n}	$\mathbf{3}$	Ψ
S_1	0	1	1	1	1	0	ϕ_1
S_2	1	1	1	1	1	1	ϕ_2
S_3	1	1	1	1	1	1	ϕ_3
S_4	1	1	1	1	1	1	ϕ_4

Table 5: Logical values of matrix S , vector Φ and the scalar Ψ calculated for population X .

Table 5 shows that Ψ for the current population against the preference information has been resolved as 3 and because $\Psi \geq T$ the Z-phase of the WZ preference articulation operator takes effect. This allows for the minimisation of the solutions within the found discovered by calculating z_{mn} for every x_{mn} and aggregating them into a single score as the scalar V_n , so that the solutions X_n may be sorted in order of descending proximity to P .

	z_{1n}	z_{2n}	z_{3n}	z_{4n}	z_{5n}	V	
Z_1	0.82	-0.71	-0.27	-0.54	0	-0.14	V_1
Z_2	-0.82	-0.44	-0.27	-0.87	-0.89	-0.66	V_2
Z_3	-1.63	-0.44	-1.63	-0.54	-1.79	-1.21	V_3
Z_4	0	-1.76	-1.09	-1.63	0	-0.9	V_4

Table 6: The matrix Z containing the z-scores for the population X , and the aggregated z-score V for each solution X_n .

In Table 6 it can be seen that z_{mn} has been calculated for every solution's objective value and has been aggregated per solution as the scalar V_n . Sorting the entries of V in ascending order resolves (as calculated by the WZ preference articulation operator) the solutions X_n within the ROI and furthest in proximity to P in objective space. In this example, the solutions are ordered (presented in order of descending proximity to P) $\langle X_3, X_4, X_2, X_1 \rangle$.

3. Benchmarking WZ-MOEA/D

In mathematics, optimisation is concerned with the selection of optimal solutions to objective functions. An objective function consists of input arguments referred to as variables (or genotype) which are computed by one or many mathematical functions to determine the objective value (or phenotype).

Real-world optimisation problems are divided into one (in the case of single objective optimisation) or many (in the case of multi-objective optimisation) objective functions in order to be optimised by an optimisation algorithm. The difficulty of convergence can be made easier by the bounding of problem variables as this reduces the size of the search domain.

In order to determine an EAs robustness, its performance must be assessed on the optimisation of objective functions created for the purpose of testing. These test function often have a scalable number of objectives and problem variables as well as a complex Pareto shape, and aim to test algorithms on their ability to converge to an approximation set in the presence of search space difficulties often present in real-world optimisation problems.

To assess the performance of the WZ preference articulation operator and whether its host algorithm benefits from any improvements, test cases have been designed where two implementations of MOEA/D-DRA (with and without the incorporation of the WZ preference articulation operator) have been executed on a range of commonly used test functions from the literature and the final approximation set has been analysed in terms of standard performance metrics. From this point onward WZ-MOEA/D will refer to an implementation of MOEA/D-DRA [40] with the incorporation of the WZ preference articulation operator.

3.1. Incorporation of the WZ preference articulation operator into MOEA/D-DRA

The WZ preference articulation operator has been incorporated into MOEA/D-DRA in order to both test the portability of the operator itself as well as the feasibility of preference articulation on test functions containing many-objectives. MOEA/D-DRA has been designed as an algorithm with the optimisation of test functions consisting of complex Pareto-sets in mind, it is important to retain the benefits of MOEA/D-DRA by incorporating the WZ preference articulation operator in a way that assists the selection process rather than completely replacing it.

WZ-MOEA/D operates in one of two phases (W-phase and Z-phase) dictated by the WZ preference articulation operator, which take effect depending on when certain criteria are satisfied, allowing the optimisation process to efficiently spend the function evaluation budget depending on the current optimisation context. Whilst the number of solutions satisfying the preference vector P is below the threshold ($\Psi < T$) the W-phase of the WZ preference articulation operator takes effect. In this phase the MOEA/D-DRA's utility selection is replaced with a selection of solutions based on their W_n score calculated using Equation 19. If during the optimisation process the threshold ($\Psi \geq T$) is satisfied then the Z-phase of the WZ preference articulation operator takes effect, whilst in this phase a modified implementation of MOEA/D-DRA's utility selection is used, where the edging sub-problems are no longer considered as elite and solutions that do not satisfy ($\phi_n = 0$) the decision maker's expressed preferences P are discarded.

Using these two phases WZ-MOEA/D is able to get close in proximity to the decision maker's expressed ROI within a small number of function evaluations, and then produce solutions within the ROI and minimise solutions whilst retaining the diversity features of MOEA/D-DRA.

3.2. Multi-objective test suites

Two test suites are considered in the comparison of WZ-MOEA/D to MOEA/D-DRA: the real-valued test problems found in the ZDT bi-objective test suite proposed in [41]; and the scalable WFG multi-objective test suite proposed in [23]. These test suites aim to incorporate combinations of features in each test problem that a MOEA may potentially

find difficult to overcome during the optimisation process, allowing assessment of a MOEAs ability in converging to the Pareto front in the presence of such difficulties.

3.2.1. ZDT test suite

The ZDT test suite contains six test functions which provide sufficient complexity to compare multi-objective optimisers: ZDT1, ZDT2, ZDT3, ZDT4, ZDT5 and ZDT6, with each function incorporating a feature that is known to cause the EMO process difficulty in convergence to the Pareto front, and the maintenance of diversity in the approximation set. Each test function has two objectives and is concerned with their minimisation.

ZDT1 is a 30 variable problem with a convex Pareto front, ZDT2 is a 30 variable problem with a non-convex Pareto front, ZDT3 is a 30 variable problem with a Pareto front consisting of non-contiguous convex parts; ZDT4 is 10 variable problem which tests the ability to handle multi-modality with 21^9 local Pareto fronts; and ZDT6 is a 10 variable problem with solutions non-uniformly distributed along the Pareto front, with the diversity of solutions decreasing near the Pareto front. ZDT5 was not included in the experiment due to the requirement for binary represented decision variables.

3.2.2. WFG test suite

The WFG test suite contains 9 test functions (WFG1-9) which are scalable to any number of parameter variables and objectives.

WFG test functions: 1; 2; 3; 6; 7 and 8 are uni-modal with the remaining being multi-modal. WFG1 has a mixed convex geometry, WFG2 has a disconnected convex geometry. WFG3 has a degenerate and linear geometry, and all remaining WFG test functions have a concave geometry.

3.3. Test cases

A test case consisting of a chosen ROI has been chosen for each test function, to test the convergence of each algorithm to different areas of objective space. These test cases have been defined in Table 7 for two-objective problems and Table 8 for five-objective problems.

3.4. Performance assessment

Both WZ-MOEA/D and MOEA/D-DRA have been executed 25 times on each test function to reduce stochastic noise. This sample size has been selected from evolutionary computation literature e.g. [7, 38, 39]. The sufficiency of this sample size was tested by producing a large number of hypervolume indicator value samples by executing WZ-MOEA/D 200 times (the distribution of which has been illustrated in Figure 15) on the WFG6 synthetic test problem using the method of performance assessment described in this section. These 200 samples were then used to identify the relationship between the standard error of the mean (SEM) and the sample size using:

$$SEM = \frac{SD}{\sqrt{N}} \quad (20)$$

Test Function	Region of Interest
ZDT1	1.0, 0.2
ZDT2	0.3, 1.0
ZDT3	1.0, -0.4
ZDT4	0.2, 1.0
ZDT6	1.0, 0.2

Table 7: Test cases defining the ROI used for each two-objective test function.

Test Function	Region of Interest
WFG1	2.5, 1.0, 1.0, 1.0, 1.0
WFG2	1.0, 0.5, 0.5, 0.5, 2.5
WFG3	0.5, 0.5, 1.0, 2.0, 10
WFG4	2.0, 1.0, 6.0, 1.0, 2.0
WFG5	5.0, 5.0, 1.5, 1.5, 1.5
WFG6	2.0, 0.5, 0.5, 1.0, 11
WFG7	0.6, 0.6, 0.6, 8.5, 0.6
WFG8	0.5, 0.5, 0.5, 9.0, 0.5
WFG9	3.0, 1.0, 1.0, 2.0, 1.5

Table 8: Test cases defining the ROI used for each five-objective test function.

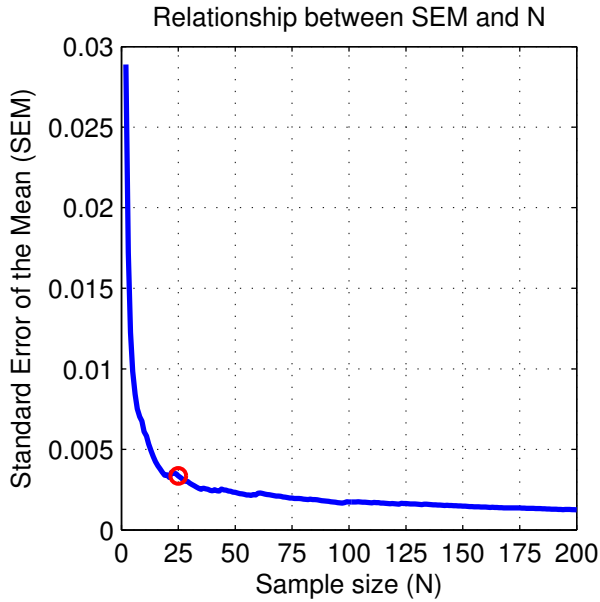


Figure 14: Relationship between standard error of the mean (SEM) and the sample size of hypervolume indicator values from 200 executions of WZ-MOEA/D on the WFG6 synthetic test problem.

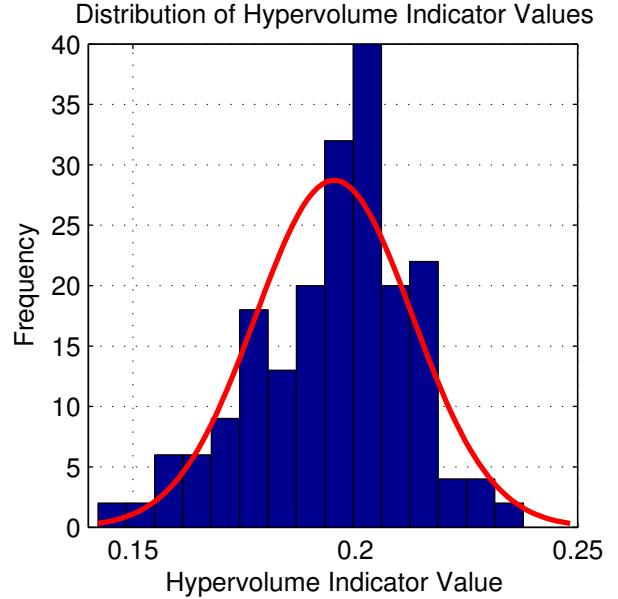


Figure 15: Histogram showing the distribution of the hypervolume indicator values from 200 executions of WZ-MOEA/D on the WFG6 synthetic test problem.

This relationship has been illustrated in Figure 14 and shows the limited benefit of more than 25 independent executions of the algorithm on the synthetic test problem.

EAs are inherently stochastic and the initial conditions that ensure the reliability of para-

metric tests cannot be satisfied [25]. In order to find the significance in contrast amongst the results obtained by WZ-MOEA/D and MOEA/D-DRA, a non-parametric test (encouraged by [10, 13]) for pairwise statistical comparison has been used. The Wilcoxon signed-ranks [37] non-parametric test (counter-part of the paired t-test) has been used with the statistical significance value ($\alpha = 0.05$), this test ranks the difference in performance of the two algorithms over each approximation set and has been used in the comparison of EAs in the literature [5, 9, 32].

The population at each generation of each execution has been scored using the hypervolume indicator, with the reference point set as the decision makers specified ROI (defined in the test case). The hypervolume indicator (or *s*-metric) is a performance metric for indicating the quality of a non-dominated approximation set, introduced by [42] and described as the “size of the space covered or size of dominated space”. It can be defined as [35]:

$$S_{f^{ref}}(X) = \Lambda \left(\bigcup_{X_n \in X} [f_1(X_n), f_1^{ref}] \times \cdots \times [f_m(X_n), f_m^{ref}] \right) \quad (21)$$

Where $S_{f^{ref}}(X)$ resolves the size of the space covered by approximation set X , $f^{ref} \in \mathbb{R}$ refers to a chosen reference point and $\Lambda(\cdot)$ refers to the Lebesgue measure. This has been illustrated in Figure 16 in two-dimensional space (to allow for easy visualisation) with a population of 3 solutions.

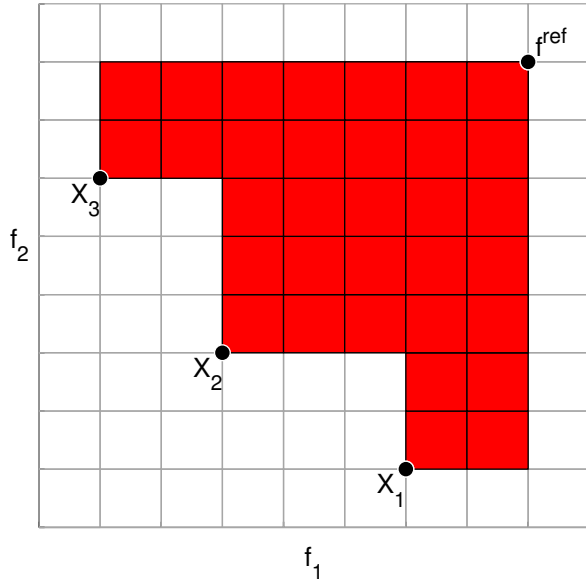


Figure 16: An example of the hypervolume indicator in two-dimensional objective space.

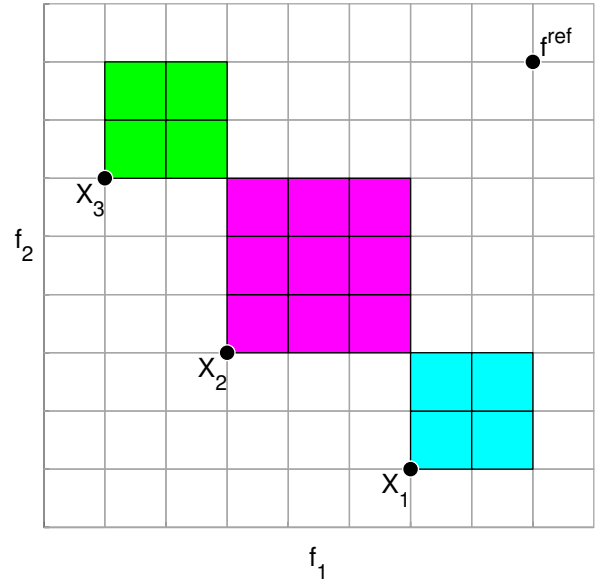


Figure 17: An example of the contributing hypervolume indicator in two-dimensional objective space.

The hypervolume indicator is selected because it is scaling independent and requires no prior knowledge of the true Pareto front, this is important when working with real-world

problems which have not yet been solved. The hypervolume indicator is currently used in the field of multi-objective optimisation as both a performance metric and in the decision making process [21, 33].

The hypervolume indicator allows performance comparison between WZ-MOEA/D and MOEA/D-DRA based on which algorithm covers the greatest amount of the search space within a specified ROI, by using the preference vector P described in Equation 7 in place of f^{ref} when calculating the hypervolume indicator value. This will be the basis for identifying which algorithm outperforms the other, an example of this measure has been given in Figure 16.

The contributing hypervolume indicator is an adaptation of the hypervolume indicator in order to be used as sorting criteria for selection operators, it has been used in the s -metric Selection Evolutionary Multi-Objective Algorithm (SMS-EMOA) [12] and the Hypervolume Estimation Algorithm for Multi-Objective Optimisation (HypE) [1]. The contributing hypervolume indicator assigns each solution in an approximation set with the size of the space covered by each solution exclusively, with this information the population can be sorted by the most dominant and diverse solutions. This has been illustrated in Figure 17 in two-dimensional space with a population of 3 solutions. The contributing hypervolume indicator is used later in Section 4 to reduce the size of the final approximation set produced by the optimiser to a size that will not overwhelm and confuse the decision maker.

3.5. Results

Plots for the results of the test cases have been illustrated in Figure 18 for two-objective test problems and in Figure 19 for five-objective test problems. Each plot presents the mean hypervolume covered within the ROI at each generation for 25 executions on a test function, for both WZ-MOEA/D and MOEA/D-DRA.

Table 9 presents information regarding the p -value resolved by the Wilcoxon signed-ranks non-parametric test for the considered synthetic test problems and a symbol indicating the observation of the null hypothesis, '+' indicates that the null hypothesis was rejected and WZ-MOEA/D displayed statistically superior performance at the 95% significance level ($\alpha = 0.05$) on the considered test synthetic test function. In all cases the null hypothesis was rejected and a statistical significance of greater than 95% was observed.

From a general observation of the results it can be seen that WZ-MOEA/D outperforms MOEA/D-DRA on each test function using the hypervolume indicator. WZ-MOEA/D succeeds in finding solutions in the specified ROI in fewer function evaluations than MOEA/D-DRA, even when MOEA/D-DRA fails to find any solutions in the ROI throughout the entire search process (within the function evaluation budget). The populations found by WZ-MOEA/D on each test function cover more of the hypervolume within the specified ROI than MOEA/D-DRA, and because of this the decision maker will have a more diverse set of candidate solutions to make a decision from, giving them a better idea of the trade-offs within the specified ROI with a better spread throughout the approximated front.

The results show that the test function posing the greatest difficulty in finding a single solution within the specified ROI is ZDT4, this can be seen in Figure 18 where the plot of the mean hypervolume indicator value over the number of generations indicates that of

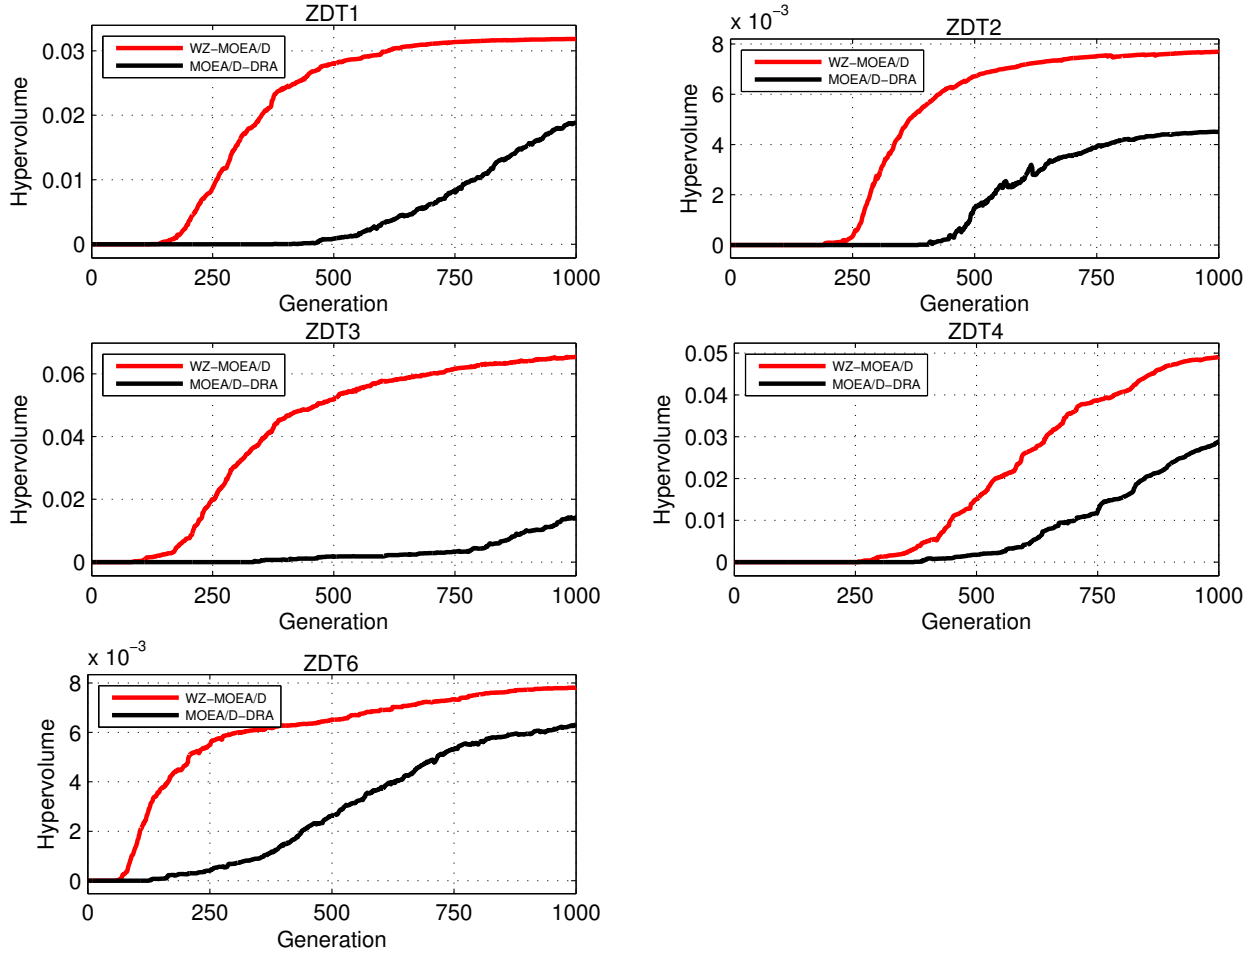


Figure 18: The mean hypervolume indicator value of WZ-MOEA/D and MOEA/D-DRA populations at each generation for bi-objective ZDT test suite.

all the test functions ZDT4 required each algorithm to search for more generations before finding solutions within the ROI. This is to be expected as a result of ZDT4’s multi-frontal nature [11] which is a struggle for most EMO algorithms; however, when comparing WZ-MOEA/D to MOEA/D-DRA it is clear that WZ-MOEA/D can get to the specified ROI faster and produce a final approximation set of better hypervolume indicator quality. In Figure 20 the populations for WZ-MOEA/D and MOEA/D-DRA have been plotted for an execution for which the hypervolume indicator value was close to the mean of each respective algorithms 25 executions, these plots show that WZ-MOEA/D’s final approximation set has converged through more of the deceptive Pareto fronts than MOEA/D-DRA, with more candidate solutions, better diversity, and also including more solutions near the extremes of the ROI. A similar result can be seen in Figure 21 where the worst of the 25 populations for WZ-MOEA/D and MOEA/D-DRA have also been plotted for ZDT4, in this case however MOEA/D-DRA has only found a single solution within the ROI where as WZ-MOEA/D has completed with better proximity and diversity.

Test Function	p -value		Test Function	p -value	
ZDT1	2.2857e-09	+	WFG3	6.132e-07	+
ZDT2	2.5742e-09	+	WFG4	0.0012886	+
ZDT3	1.1773e-07	+	WFG5	5.806e-09	+
ZDT4	1.309e-07	+	WFG6	5.8255e-05	+
ZDT6	0.00051446	+	WFG7	0.00023161	+
WFG1	7.6832e-08	+	WFG8	1.0932e-06	+
WFG2	0.022656	+	WFG9	1.7323e-06	+

Table 9: Results from pairwise comparison of the final approximation sets of both considered algorithms on each synthetic test function using the Wilcoxon signed-ranks non-parametric test.

The results for WFG6 show that both WZ-MOEA/D and MOEA/D-DRA find solutions within the ROI very early in the search, taking less than 50 generations on average. This is the result of relaxed preferences used in the test case, in particular for the 5th objective where the preference was set to 11 or below. In general, both WZ-MOEA/D and MOEA/D-DRA find the ROI in a similar number of function evaluations, with WZ-MOEA/D converging to populations with better hypervolume indicator quality very early on in the search. In the worst performing (in regards to hypervolume indicator quality) execution of each algorithm on WFG6, the plots for the final population show that WZ-MOEA/D in Figure 23 finds many solutions within the ROI with good diversity to offer the decision maker with an idea of the trade-offs, where as MOEA/D-DRA in Figure 22 fails to find a solution in the ROI.

In some of the results it can be seen that the achieved hypervolume indicator value oscillates throughout the optimisation process, for example in WFG4 and WFG7. This occurs because both WZ-MOEA/D and MOEA/D-DRA are not hypervolume indicator based algorithms and instead use a weighting based selection method, therefore the hypervolume indicator is not taken into account in any point of the optimisation process, unlike algorithms such as CMA-PAES and MO-CMA-ES, which take the contributing hypervolume indicator into account during their selection process.

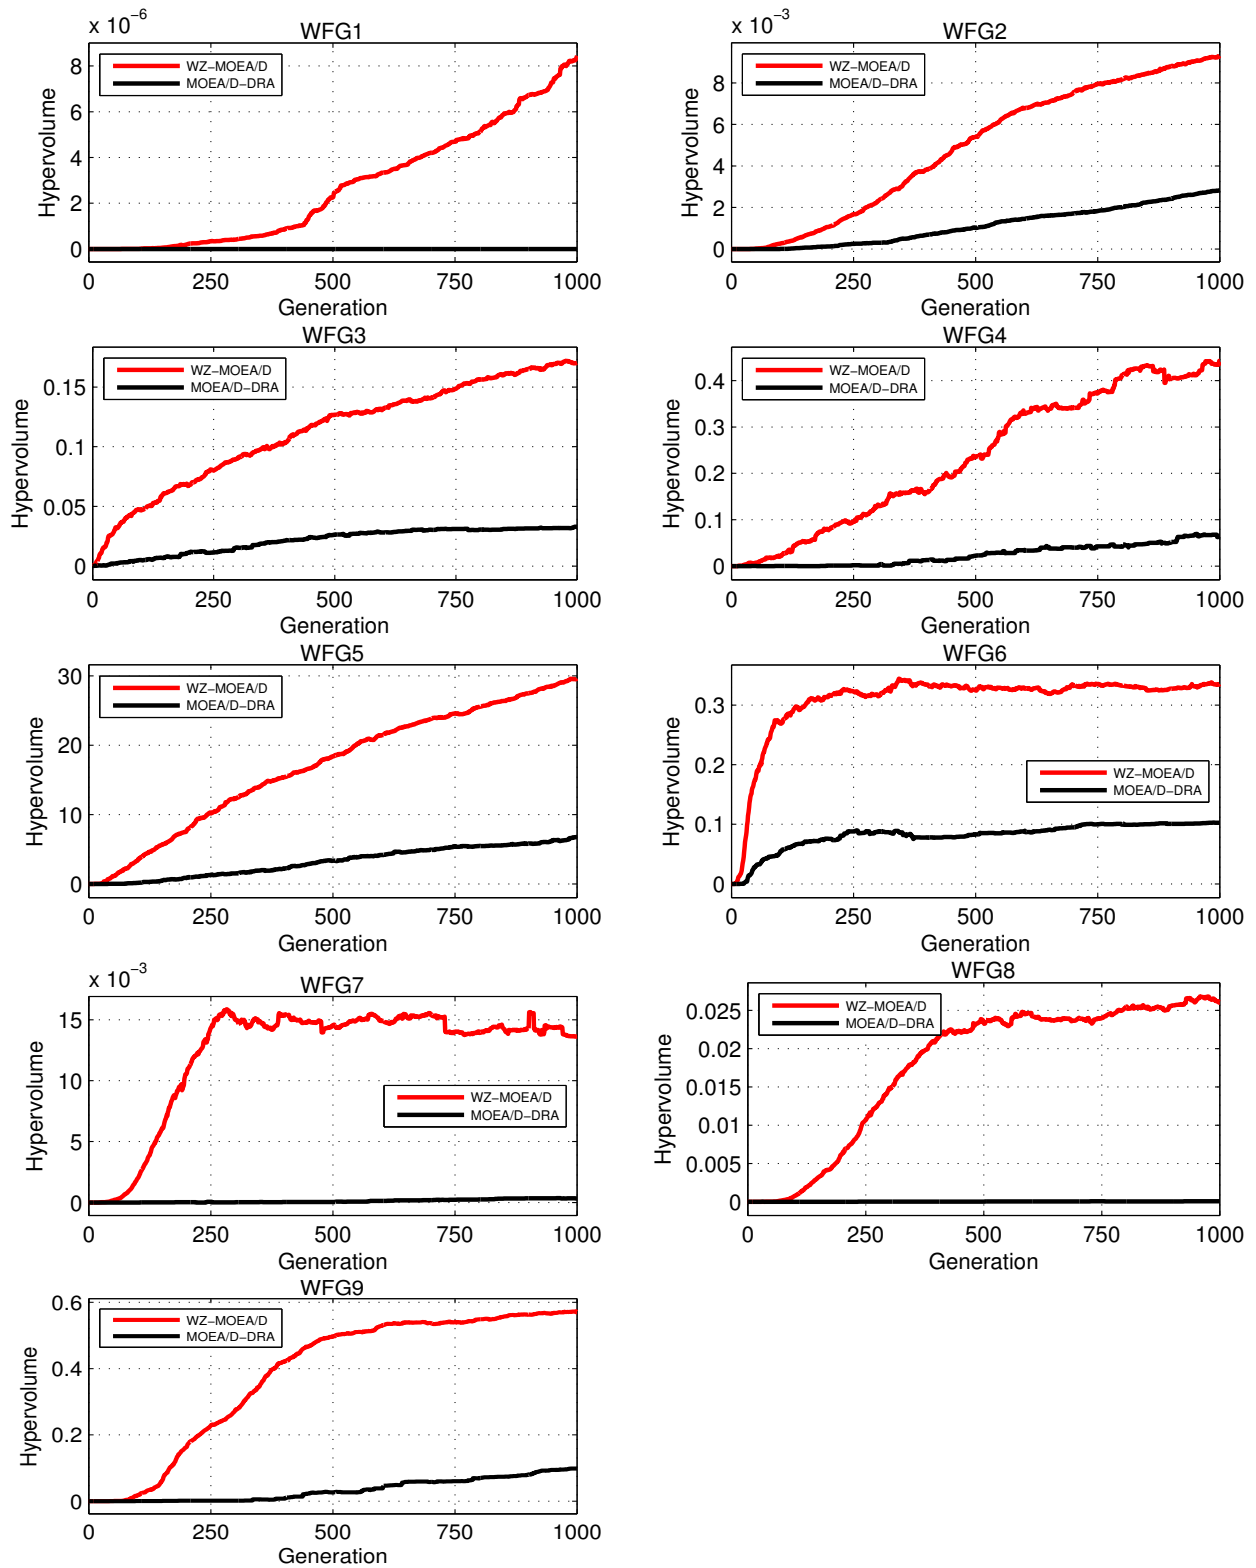


Figure 19: The mean hypervolume indicator value of WZ-MOEA/D and MOEA/D-DRA populations at each generation for five-objective WFG test suite.

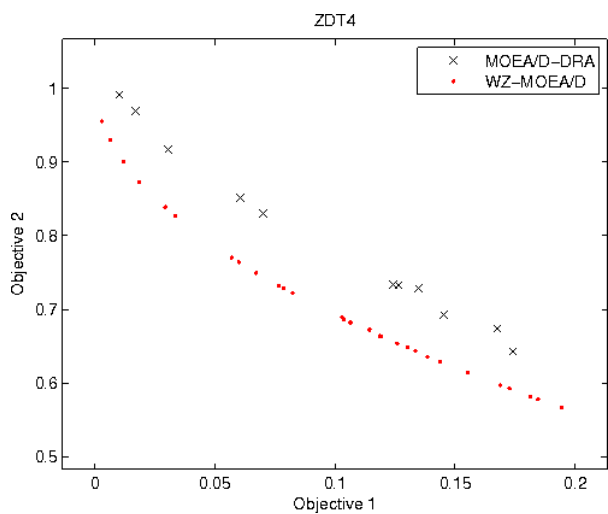


Figure 20: Population generated by worst run of WZ-MOEA/D and MOEA/D-DRA on ZDT4 within the ROI, with a hypervolume indicator value close to the mean of the 25 executions.

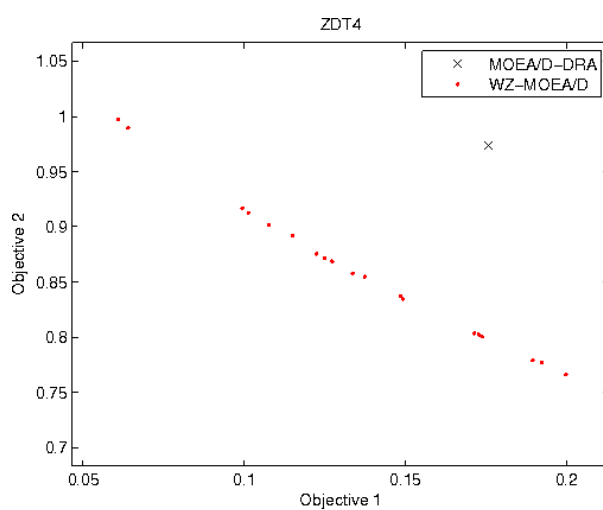


Figure 21: Population generated by worst run of WZ-MOEA/D and MOEA/D-DRA on ZDT4 within the ROI, with a hypervolume indicator value of 0.000063568 and 0.0159 respectively.

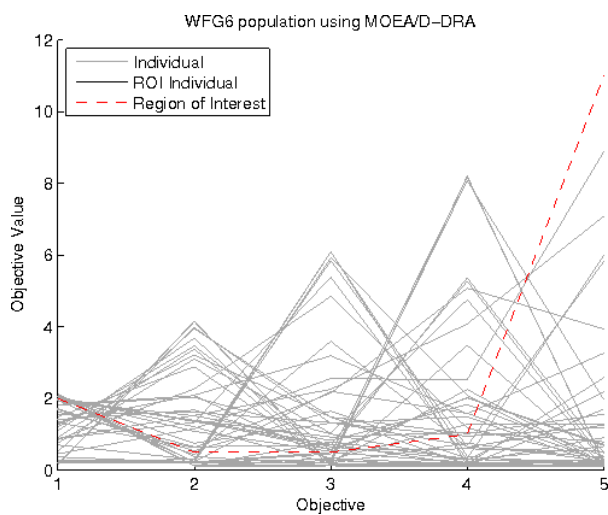


Figure 22: Population generated by worst run of MOEA/D-DRA on WFG6, with a hypervolume indicator value of 0.

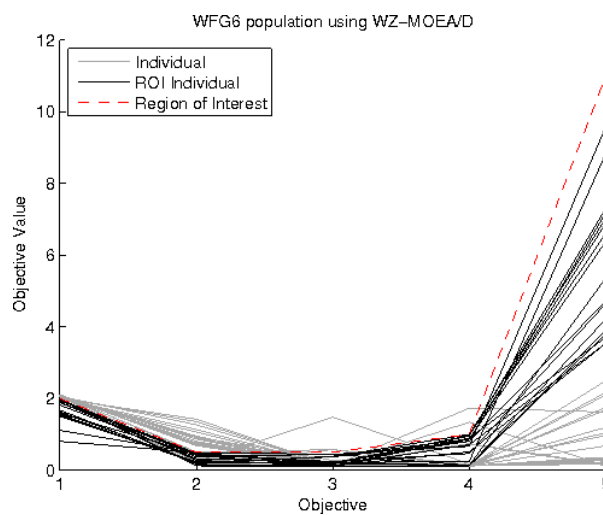


Figure 23: Population generated by worst run of WZ-MOEA/D on WFG, with a hypervolume indicator value of 0.0485.

4. Application to weapon detection

Concealed Weapon Detection (CWD) is an important area of research in the defence and security community. This is a result of a number of high profile terrorist attacks which have resulted in loss of life and damage to public infrastructure.

The threat faced by the security forces is diversifying and current technology is struggling to meet new requirements. The technologies currently in use at airports include metal detection portals, millimetre wave imaging systems, x-ray scanners and ion mobility spectrometers. These technologies are all designed to detect specific threats to security and collectively they are used to satisfy the current requirements for screening in the aviation industry. The use of multiple technologies in airports has led to choke points with lengthy waits at security checkpoints, making this method of screening undesirable and it is only used in places where it is absolutely necessary. Recently the threat of terrorism has spread and many more public areas and government buildings are becoming targets, even sporting events have now become targets. This presents a real problem as the best examples of mass screening are demonstrated at airports and as mentioned previously this type of screening leads to choke points and delays. Therefore there is a requirement for a fast method of screening people in crowded areas, that is capable of detecting a diverse range of threats.

One method of detecting a diverse of range concealed weapons in crowded areas in real-time is to use small portable radars. A number of radar systems have been developed for this purpose [3] and [27]. These radars use multiple methods of detecting concealed weapons such as time domain reflectometry and the exploitation of polarisation changes induced by complex objects concealed on the human torso under clothing.

The radar used in this work is constructed of a Vector Network Analyser (VNA) with pyramid horn antennas connected to the VNA using suitable cabling. A laptop is used to control the VNA and then classify the signals. The radar signals are analysed on the laptop using pattern recognition applied to the time resolved signals in the form of an ANN, this set-up has been illustrated in Figure 24. This method has been discussed in detail in a previous publication [28].

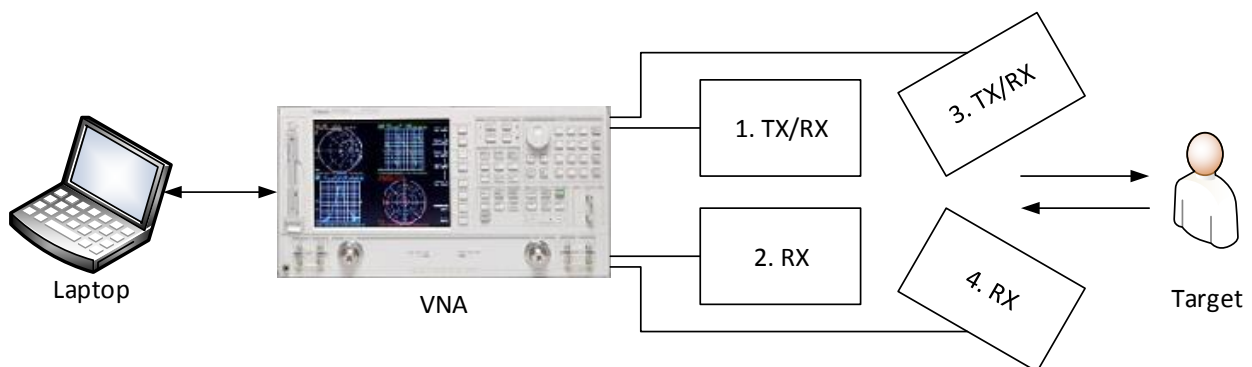


Figure 24: System block diagram illustrating the arrangement of the transmitted and receiver horn antennas.

One of the shortcomings of this method is that the optimisation of the ANN architecture has previously been performed using trial and error. This has been done by increasing the

number of hidden layers in the ANN and also increasing the number of neurons on each of these layers. A set of validation data was fed into the ANN each time a new layer or neuron was added and the true positive and false positive rates were recorded. The best architecture was selected by weighing the achieved true positive rate against the cost in false alarm rate. Another issue exists with the training of the ANN, which tends to be inconsistent. This is caused when an initial guess at the weights and biases is taken. As a result of the randomness of this guess the convergence can be to a local minima rather than the global.

It is of great importance that the false alarm rate is kept low, typically below a few percent, for security screening of large volumes of people. This is due to the action that must be taken once a potential threat has been identified. This action could range from further investigation, e.g. stop and search, to the evacuation of a crowded public area. If the false positive rate goes above a few percent the inconvenience to the security forces and general public would render the method ineffective. Therefore the primary objective in optimising the ANN architecture must be the reduction of the false positive rate.

The second objective in the optimisation of the ANN architecture must be the preservation of the true positive rate for targets of interest. The targets that should be detected by the radar include knives and guns. It is unfortunate that knives and guns are seized by the security forces far too frequently and pose a significant threat to the safety of the general public. The damage that can be caused with these weapons is considerable and these targets are easily concealed upon the human body.

4.1. Encoding the problem

In order to use evolutionary optimisation to optimise the topology and weights of the ANN classifier for concealed weapon detection, the ANNs topology and weights must be encoded into a real-valued chromosome, which can then be subjected to the various evolutionary operators used in the optimisation process and then decoded for evaluation. Figure 25 illustrates the chromosome structure used to store the problem parameters for an ANN with 2 output neurons, a maximum of 2 layers, and 8 input neurons.

HL1Neurons	HL2Neurons	ILWeights	HL1Weights
HL2Weights	HL1Bias	HL2Bias	OLBias

Figure 25: Encoded chromosome for an ANN consisting of 2 hidden layers (HL), input layer (IL), 2 neurons on the output layer (OL), and associated biases.

Parameter boundaries are also required to restrict the number of hidden layers, neurons per hidden layer, and ranges for the weights and biases within a lower and upper limit. All hidden layers but the last can contain a number of neurons ranging from none to twice the number of input neurons as seen in Equation 22, and the last hidden layer must contain a minimum of neurons equal to the number of input neurons as seen in Equation 23, this means each candidate network generated by the optimiser must have at least one hidden layer, preventing the generation of benign networks which would waste function evaluations

throughout the entire optimisation process. Finally, each weight or bias is restricted to the same boundary shown in Equation 24.

$$b(1...(HL - 1)) = \{x \in \mathbb{Z} \mid 0 \leq x \leq 2i\} \quad (22)$$

$$b(HL) = \{x \in \mathbb{Z} \mid i \leq x \leq 2i\} \quad (23)$$

$$w = \{x \in \mathbb{R} \mid -5 \leq x \leq 5\} \quad (24)$$

For the ANN configuration used in this network, each candidate solution contains 452 variables, with the first 2 defining the number of hidden layers and the number of neurons on each respectively, the following 128 variables defining the weights for the input layer, 256 for the first hidden layer, and 32 for the final hidden layer.

Regardless of the topology of the candidate solution ANN (which in this case is defined by the first 2 genes of the encoded chromosome) the maximum number of weights and biases will be stored with each chromosome, however not all genotypes will manifest and be expressed as phenotypes as only the weights and biases required to configure the candidate solutions ANN topology will be decoded and used. These unused weights and biases will remain unexpressed in the phenotype until the first two genes allow them to manifest, and can go through many generations as dormant genes, this introduces the interesting feature of atavism⁶ into this problem.

At each function evaluation a chromosome is decoded and used to create an ANN, this ANN is then used to classify the training data and the objective information is extracted based on the ANN result set.

4.2. Comparison to existing solution

To assess the performance of the optimisation algorithm when applied to this novel real world problem a benchmark must be taken. The performance of the optimised ANN is compared to previously published data [28], the experimental methodology used to obtain this data is explained here in.

The radar used in [28] is a novel multi-polarimetric frequency modulated continuous wave radar operating in the K Band (14-40 GHz). The radar illuminates a person using two linear polarisations which are switched. The use of two illuminating polarisations has been shown to significantly improve detection rates [28]. The reflected signal is then recorded in four polarisations, which are co and cross polar to each of the illuminating polarisations. The gain and phase of the reflected signals are then recorded with respect to the outgoing signal, in the frequency domain. Once collected the gain and phase information is used to generate an array of complex numbers that are then time resolved using an Inverse Fast Fourier Transform (IFFT). Examples of typical radar signals from the human body with and without concealed weapons are given in Figures 26 and 27.

The radar signals from the body with and without a concealed weapon vary with the aspect in which the person and weapon are presented in the illuminating radar beam. To

⁶“Atavism is the tendency to revert to ancestral type. In biology, an atavism is an evolutionary throwback, such as traits reappearing which had disappeared generations before.”

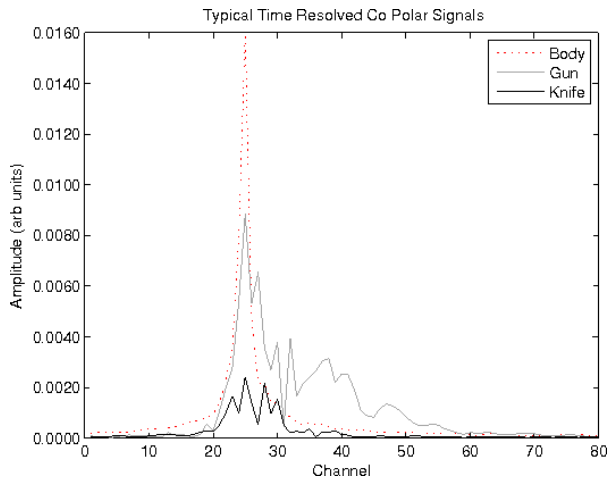


Figure 26: Typical co polarised radar signals from a body alone, a body with a concealed gun and finally a body with a concealed knife.

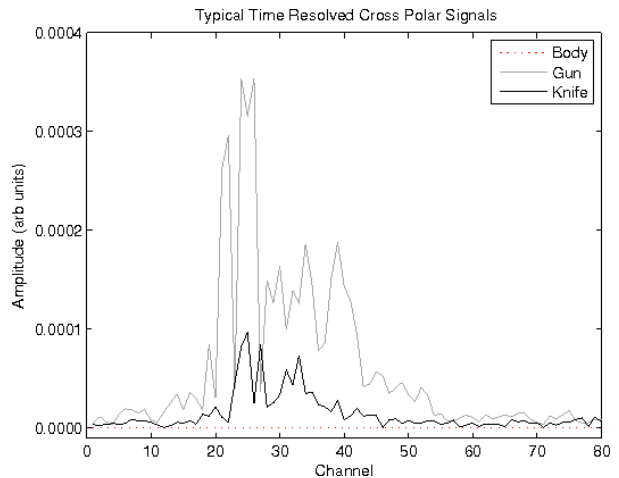


Figure 27: Typical cross polarised radar signals from a body alone, a body with a concealed gun and finally a body with a concealed knife.

compensate for this variance many scans must be taken whilst a person is moving in the beam and an accurate representation of the real operating conditions are given to the ANN. The large amounts of data obtained in the collection of training scans makes the problem difficult for the ANN to converge to a solution that is robust but without over training the ANN. In [28] a Principal Component Analysis (PCA) data reduction technique was integrated into the classification algorithm. The PCA data reduction was applied to the time resolved radar signals to obtain a set of Eigenvalues and Eigenvectors. The Eigenvalues were then used to train the ANN to classify the data into two classes, namely 'threat' and 'non-threat'. A threshold was then applied to the output neuron of the ANN and an alarm was triggered when the output value became larger than the threshold.

In [28] the architecture of the ANN is a 3 layer feed-forward network with as many input neurons as there are significant principal components (typically 5 or 6), the hidden layer has one extra neuron than the input layer and the output layer has a single neuron. The ANN was trained using a constant gradient descent back-propagation method. The training set was constructed of eigenvalues corresponding to 700 multi-polarimetric radar scans. Of these 700 scans, 100 were taken with a body without a concealed weapon and the other 600 scans with a concealed weapon. There were two different weapons used in the training set, which were a knife and a gun (300 scans of each). The validation of the ANN was performed using a dataset with the same number of scans and same distribution of body with and without a concealed weapon. The validation dataset was constructed using radar scans which the ANN had not seen *a priori*.

4.2.1. Results

WZ-MOEA/D has been used to optimise the classifier parameters for classification of concealed weapon detection with two objectives: true positives and false positives. The final

population produced by the optimiser has been plotted in Figure 28, where it can be seen that a number of trade-off solutions have been found with the ROI, resulting in a Pareto optimal approximation set. When compared to the benchmark plot it can be observed that it is Pareto dominated by every candidate solution in the final population, this means that the optimiser has found a diverse set of solutions which are all an improvement on the existing solution. The same candidate solutions have also been plotted in Figure 29 without conversion of all objectives for minimisation, this plot is what was presented to the decision maker when selection of a successful candidate solution was required.

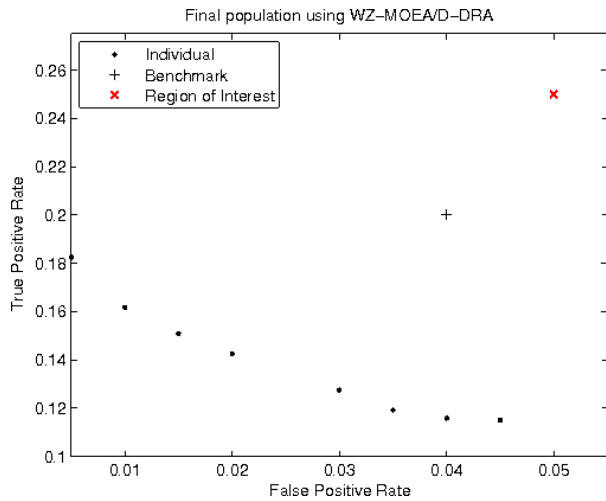


Figure 28: Population generated by MOEA/D-DRA combined with WZ preference articulation after 2000 function evaluations for the two-objective concealed weapons detection classifier.

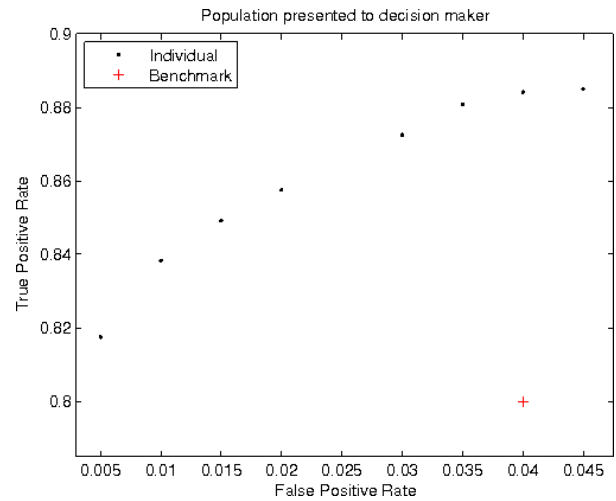


Figure 29: Population of candidate solutions presented to the decision maker, for selection to be made based on expert knowledge.

The plot in Figure 29 shows a number of optimised solutions. One of these solutions must be selected, by the decision maker, to be applied in a concealed weapon detection radar. The distribution of solutions in Figure 29 shows that each solution offers a distinctly different true positive and false positive rate. This allows the decision maker to use these solutions in a manner that will give control over the behaviour of the weapon detector. This will enable the decision maker to choose a solution that increases the system's sensitivity, when there is a heightened threat to security. Equally the false alarm rate could be reduced by choosing another solution.

The benchmark shown in Figure 29 was taken from [28], this solution offers a true positive rate of 0.8 at a cost in false positive rate of 0.04. Given the requirement for false alarm reduction which was identified as the primary objective earlier the solution from Figure 29 that was chosen, by the decision maker, was the solution offering a true positive rate of 0.82 and a false positive rate of 0.005. The chosen solution has provided a false alarm reduction of 0.035 at the same time as increasing the true positive rate by 0.02. To put this into perspective previously 40 people in one thousand would have been wrongly accused of carrying a concealed weapon whereas the number of wrongly accused with the

new optimised solution would be less than five in a thousand. Not only does this reduce the inconvenience to the security forces and general public it builds confidence in the ability and robustness of the weapons detection system.

4.3. Detection and classification of multiple types of threat

A weapon detection system that is capable of classifying a detected threat into target groups would be an extremely valuable tool to security forces. Such a system would enable the authorities to react to a detection in a controlled and proportional manner. The action that must be taken to confront an individual concealing a threat object depends very much on the threat object itself. An individual carrying a knife could be dealt with easier than an individual with an improvised explosive device.

The extent of the ROI will be determined by the decision maker based on some pre-determined criteria, for example the radar may be deployed in an environment where the client has specified a minimum detection rate and maximum false alarm rate. It is possible that no solutions may be found within a ROI which has been confined based on a clients specifications, in which case the specification would be deemed beyond the performance of the radar and another solution would be required.

As the number of targets of interest increases so does the risk of increasing the number of conflicting objectives. The extent of which this method can handle conflicting objectives will be explored in more detail in subsequent publications.

It should be noted that the radar beam is confined to a spot size which is commensurate in size with the weapons of interest, therefore it is infeasible to measure more than one weapon in a given location on the body. The authors suggest that the radar operator would scan the radar beam over the person and would be able to find multiple weapons concealed on different parts of the body and then address the situation based on the most severe threat detected.

To investigate whether the developed radar based weapon detection system is capable of classifying the reflected signal into target groups an experiment has been conducted. In this experiment radar signals from a body without a concealed weapon, a body with a concealed knife and a body with a concealed gun have been recorded by the radar. Again the radar signals were recorded in the frequency domain with both amplitude and phase relative to the illuminating beam measured. The amplitude and phase were used to generate an array of complex numbers that were temporally resolved using an IFFT, the resulting signal was then reduced using PCA as described earlier. This time the training set was constructed of 300 sets of Eigenvalues correspond to the radar scans, 100 scans of each target scenario were taken. That is 100 scans from body alone, 100 with a concealed knife and 100 with a concealed gun. The validation set was constructed of the same number of scans and the same distribution of targets. The scans used to construct the validation set had not been seen by the ANN *a priori*. The architecture, in terms of the number of layers and number of neurons on each layer, of the ANN was determined by the optimisation algorithm and the weights and biases were also determined by the optimisation algorithm.

To assess the statistical significance of these results the area under the observed ROC curve (as observed in [28]) was used alongside the chosen sample sizes to calculate the

standard error.

$$\begin{aligned}\theta_1 &= 0.8362 \\ n_A &= 200 \\ n_N &= 100\end{aligned}\tag{25}$$

θ_1 is the area under the curve for the observed ROC, n_A is the number of scans with a weapon and n_N is the number of scans without a weapon. The standard error is calculated as follows:

$$SE = \sqrt{\frac{\theta_1(1 - \theta_1) + (n_A - 1)(Q_1 - \theta_1^2) + (n_N - 1)(Q_2 - \theta_1^2)}{n_A n_N}}\tag{26}$$

Where Q_1 and Q_2 are calculated using the following equations

$$Q_1 = \frac{\theta_1}{2 - \theta_1}\tag{27}$$

$$Q_2 = \frac{2\theta_1^2}{1 + \theta_1}\tag{28}$$

The standard error was calculated to be 0.0225, this shows a high level of certainty that the results are statistically significant. It can be shown through testing a null hypothesis (that the results happened by chance) that the sample size chosen gives a high level of confidence in the significance of these results. To do this the area under the curve is set at 0.5 this represents a special case of a ROC which is by chance and has no useful classification abilities. The difference between the area under the observed ROC and the area under the curve for the null hypothesis is divided by the standard error to give a z-score. A z-score of 1.645 relates to a 5% one-sided test of significance and a 95% power.

$$z = \frac{\theta_1 - \theta_2}{SE}\tag{29}$$

$$\theta_2 = 0.5\tag{30}$$

The value of the z-score calculated for this test was 14.9, this shows that the sample sizes chosen were much larger than required and thus result in a high level of statistical significance.

4.3.1. Results

The population of solutions from the WZ-MOEA/D optimisation are plotted in Figure 30, this plot shows the solutions for a 5 objective problem. The 5 objectives are split into true positives for the body, gun and knife and the false positives for the gun and knife. In this problem the reduction of both false positive rates are the two main objectives, these should be weighted equally. The remaining objectives are the maximization of the true positive rates for each target included in the training and validation sets. Also presented is

a colour-map, Figure 31, this was used to aid the decision maker in choosing a solution from the population. In each of the presented plots there are 5 candidate solutions, one of which should be chosen by the decision maker to be implemented in the weapon detection system.

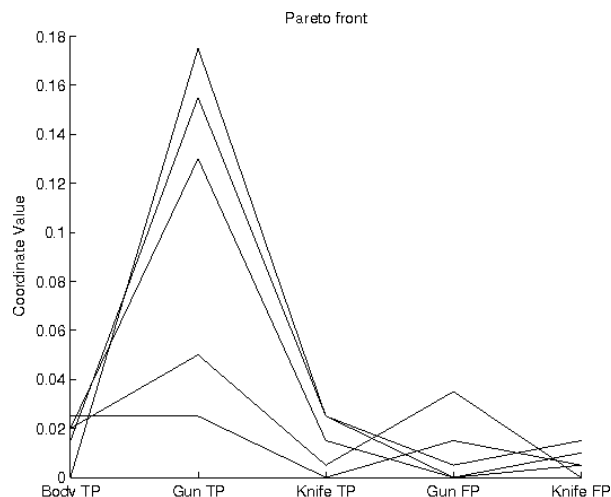


Figure 30: Objective results for 5-objective threat detection.

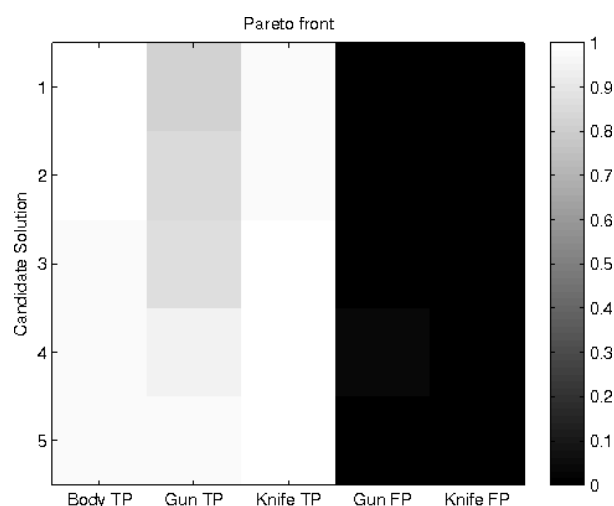


Figure 31: Colour-map used by decision maker.

The colour-map was found to be useful tool in the selection of a solution from the final candidate population. When a population of 5 solutions was presented to the decision maker in the minimisation format presented in Figure 30 it was difficult to visualise the different trade-offs between candidate solutions. The visualisation of the trade-off between solutions is much clearer in the colour-map in Figure 31. Figure 31 shows that all solutions provide a means to reliably classify the signals into targets classes which are body, gun and knife; where the true positive rates are high and the false positive rates are low for all target classes. This can also be seen in the minimisation format presented in Figure 30. In Figure 31 it is clear to see that there are subtle differences in the trade-offs between the 5 objectives in the final candidate population. As the false positive rates for both target classes are effectively equal for all candidate solutions, the solution taken by the decision maker was chosen as the solution with near equal true positive rate for all targets. This solution is candidate solution number 5.

5. Conclusions

In this paper a novel method of preference articulation has been introduced in the form of the WZ preference articulation operator. The WZ preference articulation operator has been incorporated into an implementation of MOEA/D and benchmarked on bi-objective and five-objective test functions and compared to MOEA/D-DRA. The results show that the WZ preference articulation operator has successfully improved the performance of its host algorithm (MOEA/D) when searching towards an expressed ROI, by producing more

solutions within the ROI, producing solutions in the ROI within fewer function evaluations, and also producing populations of better hypervolume indicator quality.

The application of the WZ-MOEA/D optimisation algorithm to the training and optimisation of the topology of an ANN intended for use in Concealed Weapon Detection has been presented. The performance of the optimised ANN has been benchmarked against previously published data, in which an ANN had been trained using back-propagation and its topology determined by trial and error. The optimisation has been shown to provide the decision maker with a number of solutions (trained ANNs) that all have independent trade-offs which are equally distributed across the Pareto front. The decision maker then selected an optimised solution which provided a reduction in false alarm rate and an increase in detection rate.

Although the results have shown that the weapons detection radar can achieve high detection rates with very little cost in terms of false alarm rates it should be noted that these values only apply in the scenario used to collect the data published within in this paper. It is anticipated that the system performance will fall away and fewer detections will be made in other scenarios, for example if the person being screened is non-co-operative then the operator may be unable to make a full scan of the person and therefore a concealed weapon may go undetected. The application of radar to weapon detection is novel and therefore all operational procedures and system performance have not yet been evaluated, this will be subject to further research.

There is an opportunity for further work in using the WZ preference articulation operator in combination with an offline archive to explore the effects of changing preferences during the optimisation process and using the offline archive to search for solutions matching the new preferences, to identify whether this type of combination would speed up search in the presence of shifting preferences, without the need for re-starts. There is also an opportunity for further work with the development of a multi-objective test suite targeting the performance assessment of the ability for an optimiser to incorporate decision maker preferences, this would also provide a standard for future comparison and assessment of preference driven optimisers.

References

- [1] J. Bader, E. Zitzler, Hype: An algorithm for fast hypervolume-based many-objective optimization, *Evolutionary Computation* 19 (1) (2011) 45–76.
- [2] E. Besada-Portas, L. De La Torre, A. Moreno, J. L. Risco-Martín, On the performance comparison of multi-objective evolutionary uav path planners, *Information Sciences* 238 (2013) 111–125.
- [3] N. Bowring, M. Southgate, D. Andrews, N. D. Rezgui, S. Harmer, D. O’Reilly, Development of a longer range standoff millimetre wave radar concealed threat detector, in: *Proc. SPIE 8714, Radar Sensor Technology XVII*, SPIE, 2013, pp. 1–8.
- [4] S. Chaudhuri, K. Deb, An interactive evolutionary multi-objective optimization and decision making procedure, *Applied Soft Computing* 10 (2) (2010) 496–511.
- [5] P. Civicioglu, Backtracking search optimization algorithm for numerical optimization problems, *Applied Mathematics and Computation* 219 (15) (2013) 8121–8144.
- [6] C. A. Coello Coello, Research directions in evolutionary multi-objective optimization, *Evolutionary Computation Journal* 3 (3) (2012) 110–121.

- [7] I. De Falco, A. Della Cioppa, D. Maisto, U. Scafuri, E. Tarantino, Biological invasion–inspired migration in distributed evolutionary algorithms, *Information Sciences* 207 (2012) 50–65.
- [8] K. Deb, *Multi-Objective Optimization using Evolutionary Algorithms*, John Wiley and Sons, New York, 2001.
- [9] J. Derrac, C. Cornelis, S. García, F. Herrera, Enhancing evolutionary instance selection algorithms by means of fuzzy rough set based feature selection, *Information Sciences* 186 (1) (2012) 73–92.
- [10] J. Derrac, S. García, D. Molina, F. Herrera, A practical tutorial on the use of nonparametric statistical tests as a methodology for comparing evolutionary and swarm intelligence algorithms, *Swarm and Evolutionary Computation* 1 (1) (2011) 3–18.
- [11] J. J. Durillo, A. J. Nebro, C. Coello Coello, F. Luna, E. Alba, A comparative study of the effect of parameter scalability in multi-objective metaheuristics, in: *Evolutionary Computation, 2008. CEC 2008. (IEEE World Congress on Computational Intelligence)*. IEEE Congress on, IEEE, 2008, pp. 1893–1900.
- [12] M. Emmerich, N. Beume, B. Naujoks, An emo algorithm using the hypervolume measure as selection criterion, in: *Evolutionary Multi-Criterion Optimization*, Springer, 2005, pp. 62–76.
- [13] M. Epitropakis, V. Plagianakos, M. Vrahatis, Evolving cognitive and social experience in particle swarm optimization through differential evolution: A hybrid approach, *Information Sciences: an International Journal* 216 (2012) 50–92.
- [14] M. Farina, P. Amato, On the optimal solution definition for many-criteria optimization problems, in: J. Keller, O. Nasraoui (eds.), *Proceedings of the NAFIPS-FLINT International Conference, 2002*, pp. 233 – 238.
- [15] P. J. Fleming, R. C. Purshouse, R. J. Lygoe, Many objective optimization: An engineering perspective, in: C. A. C. Coello, A. H. Aguirre, E. Zitzler (eds.), *Proceedings of the International Conference on Evolutionary Multi-Objective Optimization (EMO2005)*, vol. 3470 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin, 2005, pp. 14 – 32.
- [16] C. M. Fonseca, P. J. Fleming, Genetic algorithms for multi-objective optimization: Formulation, discussion and generalization, in: S. Forrest (ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms*, Morgan Kaufmann, 1993, pp. 416–423.
- [17] C. M. Fonseca, P. J. Fleming, Multiobjective optimization and multiple constraint handling with evolutionary algorithms - Part I: A unified formulation, *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 28 (1) (1998) 26 – 37.
- [18] F. W. Gembicki, *Vector optimization for control with performance and parameter sensitive indices*, Ph.D. thesis, Case Western Reserve University, Cleveland, Ohio (1974).
- [19] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, MA., 1989.
- [20] I. A. Griffin, P. Schroder, A. J. Chipperfield, P. J. Fleming, Multi-objective optimization approach to the alstom gasifier problem, *Proceedings of the institute of mechanical engineers* 214 (1) (2000) 453–468.
- [21] M. Helbig, A. P. Engelbrecht, Performance measures for dynamic multi-objective optimisation algorithms, *Information Sciences* 250 (2013) 61–81.
- [22] G. Howard, E. Gale, L. Bull, B. de Lacy Costello, A. Adamatzky, Evolution of plastic learning in spiking networks via memristive connections, *Evolutionary Computation, IEEE Transactions on* 16 (5) (2012) 711–729.
- [23] S. Huband, P. Hingston, L. Barone, L. While, A review of multiobjective test problems and a scalable test problem toolkit, *Evolutionary Computation, IEEE Transactions on* 10 (5) (2006) 477–506.
- [24] C.-L. Hwang, A. S. M. Masud, *Multiple Objective Decision Making - Methods and Applications*, vol. 164 of *Lecture Notes in Economics and Mathematical Systems*, Springer-Verlag, Berlin, 1979.
- [25] K. Li, S. Kwong, J. Cao, M. Li, J. Zheng, R. Shen, Achieving balance between proximity and diversity in multi-objective evolutionary algorithm, *Information Sciences* 182 (1) (2012) 220–242.
- [26] Z. Michalewicz, D. B. Fogel, *How to Solve It: Modern Heuristics*, Springer, Berlin, 2000.
- [27] D. O’Reilly, N. Bowring, S. Harmer, Signal processing techniques for concealed weapon detection by

- use of neural networks, in: IEEE 27th Convention of Electrical and Electronic Engineers in Israel IEEE, IEEE, 2012, pp. 1–4.
- [28] D. O'Reilly, N. Bowring, N. D. Rezgui, D. Andrews, S. Harmer, Remote concealed threat detection by novel classification algorithms applied to multipolarimetric uwb radar, in: Proc. SPIE 8714, Radar Sensor Technology XVII, SPIE, 2013, pp. 1–8.
- [29] R. C. Purshouse, On the evolutionary optimisation of many objectives, Ph.D. thesis, Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, UK, S1 3JD (2003).
- [30] S. Rostami, A. Shenfield, Cma-paes: Pareto archived evolution strategy using covariance matrix adaptation for multi-objective optimisation, in: Computational Intelligence (UKCI), 2012, IEEE, 2012, pp. 1–8.
- [31] A. Shenfield, P. J. Fleming, M. Alkarouri, Computational steering of a multi-objective evolutionary algorithm for engineering design, Engineering Applications of Artificial Intelligence 20 (8) (2007) 1047–1057.
- [32] B. Trawiński, M. Smetek, Z. Telec, T. Lasota, Nonparametric statistical analysis for multiple comparison of machine learning regression algorithms, International Journal of Applied Mathematics and Computer Science 22 (4) (2012) 867–881.
- [33] K. Van Moffaert, M. M. Dragan, A. Nowé, Hypervolume-based multi-objective reinforcement learning, in: Evolutionary Multi-Criterion Optimization, Springer, 2013, pp. 352–366.
- [34] C. Vlachos, D. Williams, J. Gomm, Genetic approach to decentralised pi controller tuning for multi-variable processes, IEE Proceedings-Control Theory and Applications 146 (1) (1999) 58–64.
- [35] T. Voß, N. Hansen, C. Igel, Improved step size adaptation for the mo-cma-es, in: Proceedings of the 12th annual conference on Genetic and evolutionary computation, ACM, 2010, pp. 487–494.
- [36] Y. Wang, B. Li, T. Weise, Two-stage ensemble memetic algorithm: Function optimization and digital iir filter design, Information Sciences 220 (2013) 408–424.
- [37] F. Wilcoxon, Individual comparisons by ranking methods, Biometrics bulletin (1945) 80–83.
- [38] Z. Yang, K. Tang, X. Yao, Large scale evolutionary optimization using cooperative coevolution, Information Sciences 178 (15) (2008) 2985–2999.
- [39] A. Zamuda, J. Brest, B. Boskovic, V. Zumer, Differential evolution for multiobjective optimization with self adaptation., in: IEEE Congress on Evolutionary Computation, 2007, pp. 3617–3624.
- [40] Q. Zhang, W. Liu, H. Li, The performance of a new version of moea/d on cec09 unconstrained mop test instances, in: Evolutionary Computation, 2009. CEC'09. IEEE Congress on, IEEE, 2009, pp. 203–208.
- [41] E. Zitzler, K. Deb, L. Thiele, Comparison of multiobjective evolutionary algorithms: Empirical results, Evolutionary computation 8 (2) (2000) 173–195.
- [42] E. Zitzler, L. Thiele, An evolutionary algorithm for multiobjective optimization: The strength pareto approach, Citeseer, 1998.
- [43] E. Zitzler, L. Thiele, M. Laumanns, C. M., V. G. da Fonseca, Performance assessment of multiobjective optimizers: An analysis and review, IEEE Transactions on Evolutionary Computation 7 (2) (2003) 117–132.