

ParEGO: A Hybrid Algorithm with On-line Landscape Approximation for Expensive Multiobjective Optimization Problems

Joshua Knowles

Abstract— This paper concerns multiobjective optimization in scenarios where each solution evaluation is financially and/or temporally expensive. We make use of nine relatively low-dimensional, non-pathological, real-valued functions, such as arise in many applications, and assess the performance of two algorithms after just 100 and 250 (or 260) function evaluations. The results show that NSGA-II, a popular MOEA, performs well compared to random search, even within the restricted number of evaluations used. A significantly better performance (particularly in the worst case) is, however, achieved on our test set by an algorithm proposed herein — ParEGO — which is an extension of the single-objective *efficient global optimization* (EGO) algorithm of Jones et al. ParEGO uses a design-of-experiments inspired initialization procedure and learns a Gaussian processes model of the search landscape, which is updated after every function evaluation. Overall, ParEGO exhibits a promising performance for multiobjective optimization problems where evaluations are expensive or otherwise restricted in number.

Keywords: multiobjective optimization, expensive black-box functions, EGO, DACE, NSGA-II, landscape approximation, Kriging, metamodels, response surfaces, Pareto optima, test suites, performance assessment

I. INTRODUCTION

As scientists and engineers become increasingly aware of the potential benefits of applying heuristic search methods in their disciplines, a new and greater variety of real-world optimization problems is gradually coming to light. Many of the problems fall right into the ‘backyard’ of evolutionary computation (EC) methods — especially those for which conventional techniques are not easily adapted, including non-convex, mixed integer, non-linear, constrained and/or noisy cost functions; for all these problems we may happily look forward to a wealth of EC success stories. But some other problems now appearing seem to pose a particular challenge to evolutionary computation and other heuristic search methods because they combine the requirement of optimizing multiple, incommensurable objectives with the feature of being prohibitively expensive to evaluate. It is on these latter problems that we focus in this paper.

More precisely, our motivation comes from a number of optimization scenarios arising in the experimental sciences. These have a more-or-less common list of features that we can take to loosely specify a class of problems, as follows:

- 1) the problem has multiple, possibly incommensurable, objectives,

- 2) the time taken to perform one evaluation is of the order of minutes or hours,
- 3) only one evaluation can be performed at one time (no parallelism is possible),
- 4) the total number of evaluations to be performed is limited by financial, time or resource constraints,
- 5) no realistic simulator or other method of approximating the full evaluation is readily available,
- 6) noise is low (repeated evaluations yield very similar results),
- 7) the overall gains in quality (or reductions in cost) that can be achieved are high,
- 8) the search landscape is locally smooth but multimodal,
- 9) the dimensionality of the search space is low-to-medium.

Although problems exhibiting (many of) these features include some familiar ones from engineering design optimization (e.g. [29]), there are also less well-known combinatorial biochemistry and materials science applications [14], [24], [58], as well as instrument configuration problems [54], [55], in this class.

An example of the latter which motivates the work herein is [45]. It reports on a series of experiments directed at improving the viability of *metabolomics*¹, by optimizing the configuration of a GC-TOF mass spectrometer (GC-MS). Three objectives in the optimization were considered: (i) maximizing the signal-to-noise ratio in the chromatogram; (ii) maximizing the number of ‘true’ peaks in the chromatogram, and (iii) minimizing the processing time — the time for the instrument to analyse one sample. Features 2–4 in the list given above arise in this particular problem because the evaluations are actually ‘wet experiments’² that must be performed on a physical machine (which is itself very costly and hence only one is available), and may also require costly consumables and/or the use of highly-skilled operators. Feature 5 applies here because it is practically infeasible to simulate accurately the output of such a complicated instrument over a range of chemical inputs. Thus, considering 2-5, it is clear that the EA or other optimization method *must* search the space in a smaller number of evaluations than in many other EA applications.

Fortunately, features 6–9 reduce the difficulty of the task

¹Metabolomics [2] relies on efficiently identifying the hundreds or thousands of different compounds within a biological sample to enable metabolic processes to be accurately monitored.

²The term used in biology for laboratory as opposed to computer-based (‘dry’) experiments.

somewhat. Low noise means that individual experiments need not be repeated, and there were only eight real-valued settings of the GC-MS (i.e. search space dimensions). Moreover, it is known from human experts that getting the configuration ‘right’ yields significant improvement in the chromatograms; since this is usually achieved from hand-tuning by experts the search space must be, at least locally, smooth. Therefore, reasonably large basins of attraction probably exist. However, it is also known that some limited epistasis does exist [54].

In the experiments reported in [45], a multiobjective evolutionary algorithm (MOEA), PESA-II [12], was used to perform the optimization [45] and some 180 GC-MS configurations were assayed in all (taking several days and consuming expensive instrument-time). At the end, a particular configuration that offered the best compromise solution was chosen from the Pareto front obtained. This yielded a threefold increase in the number of peaks observed, coupled with a small increase in the potential throughput (i.e., future samples could be processed more quickly), over the hand-tuned configuration usually employed. The choice of employing an MOEA was motivated primarily by the feature 1 above, and the wish to explore the tradeoffs of the three objectives, enabling the input of expert knowledge in the final choice of GC-MS configuration. The ability to do this from a single optimization run, with little or no knowledge of the cost landscape, makes an MOEA approach a logical choice. Moreover, results reported in [45] show that GC-MS configurations from later experiments were statistically significantly better than earlier configurations (computed using a Fisher’s permutation test), suggesting that the MOEA used was better than a random search would have been.

Nonetheless, the scientific justification for the use of MOEAs in this particular optimization context is not as solid as we would like because, although MOEAs have been applied in a wealth of different contexts and have exhibited competitive performance against several other optimization techniques, their performance when the number of function evaluations is severely limited is largely untested in the literature (though see [36] for some experiments underpinning the use of PESA-II for the GC-MS optimization). A *central aim of this paper* is thus to undertake testing of an MOEA in contexts similar to the GC-MS configuration problem, over a *realistically small number* of function evaluations. To this end, we choose to employ NSGA-II, a well-respected, modern example of an MOEA³, and test it, using a small population size, over just 100 and 260 function evaluations.

In many engineering applications that share similar features to our scenario above, standard evolutionary algorithms are eschewed because of their low efficiency with respect to function evaluations. A wealth of other global optimization methods (see next section) use more principled methods of exploring the search space under these restricted conditions, most notably those methods based on modeling the landscape on-line during the search. However, to our knowledge, the performance of these methods has never been compared with

an MOEA in a straightforward Pareto optimization context. Thus, a *second aim of this paper* is to extend such an algorithm to do multiobjective optimization, and to compare it with the NSGA-II. We choose for this, a frequently cited algorithm from the global optimization literature, specifically designed for expensive functions of low dimension — the *efficient global optimization* (EGO) algorithm [35]. The version that we propose for Pareto optimization is given the designation ‘ParEGO’ from here on.

To undertake the performance assessment of ParEGO and NSGA-II, we do not use expensive functions here, but employ a suite of test functions enabling us to collect performance data over multiple algorithm runs. The selection of problems in the suite includes functions exhibiting a diverse range of problem difficulties, yet is restricted to problems that we think model scientific/engineering contexts like the GC-MS optimization problem. For this reason, the problems are relatively low-dimensional (having up to 8 real-valued dimensions), are not pathologically rugged, and have either two or three objectives. In measuring performance, we give some rough judgments on the worst-case performance in addition to applying normal statistical tests to estimate the location of the distribution. Worst-case performance is of particular relevance in scenarios where only one shot at the experiments/optimization is usually possible, although it is difficult to estimate empirically. The choice of test functions and the methods for measuring performance used here should provide a basis for future comparison of other algorithms for expensive, low dimensional multiobjective problems. To facilitate this, problems, assessment methods and results are available for download at [37].

Organization

The rest of the paper is organized as follows. In Section 2, a review of relevant literature in engineering optimization methods, landscape approximation, and MOEAs, is given. An outline of EGO, a powerful global optimization method for expensive functions, is provided in Section 3, together with results demonstrating our implementation of the approach. Section 4 extends EGO to the multiobjective algorithm, ParEGO, which we propose here. Section 5 describes the choice of test functions, and Section 6 details our methods for performance assessment and comparison. Section 7 sets out the method of comparison, giving all parameter settings, Section 8 presents the results, and Section 9 provides a summary and conclusion.

II. RELATED WORK

The rise of evolutionary algorithms for multiobjective optimization in recent years is now well documented, with thorough reviews of the history and current state of the art available in [6], [9], [16]. Much of the success that the field is enjoying relies on the flexibility with which MOEAs can deal with various optimization contexts and features, such as: high dimensionality; integer, real and mixed parameters; lack of knowledge about ranges of fitness functions; non-differentiability of the cost landscape; constraints, and so on. The initial high computational overhead of some early methods has also been overcome by a gradual progression with the

³The IEEE TEC paper describing NSGA-II for multi-objective optimization was judged as the ‘Fast-breaking Paper in Engineering’ by Web of Science (ESI) in February 2004.

introduction of more efficient algorithms such as NPGA [30] and PAES [40], which first attempted to use reduce the cost of niching, through to the Micro-GA [7], the popular NSGA-II [17], and, more recently, particle-swarm optimization algorithms [8]. More efficient data-structures that can be used to speed up the niching, selection and archiving processes that these MOEAs rely upon are also now available [32].

In engineering, MOEAs have found many applications [15], from airplane wing optimization [29], [44] and spacecraft trajectory planning [13] to the design of irrigation networks [5]; several of these obviously entail expensive simulation or experimental steps. Nonetheless, and perhaps due to the focus made on improvements to the computational overhead of MOEAs *per evaluation*, their effectiveness in applications where very few function evaluations can be afforded has only recently been investigated in a few isolated papers. For example, two MOEAs that sample the search space by making explicit use of all previously visited points, and explicitly avoid, oversampling of fit regions, were proposed in [31]. And in [25] the case of dynamic multiobjective optimization, where the fitness function changes over time, requiring a rapid re-optimization to track it, was tackled using a hybrid approach combining deterministic local search routines that are efficient in function evaluations, with a (1+1)-ES. These kinds of methods appear to be promising, outperforming more conventional MOEAs on the test problems used, when the number of function evaluations is the limiting factor. Nonetheless, it would seem that further savings of precious function evaluations could be made by *learning* the cost landscape from all previously visited points, and by using this to estimate the ‘best’ place to sample next (either to improve the model the most or to obtain a better solution, or both).

Learning a cost landscape from a set of solution/cost pairs is variously called surrogate, approximate or meta-modeling in the literature. The idea of selecting the next solution to sample in order to maximize the improvement in the model is known as *active learning* [10] in the machine intelligence literature. These approaches have received little attention (and only very recently) in evolutionary multiobjective optimization but meta-modeling has a relatively rich history in optimization of a single objective. In design engineering, meta-modeling is usually known as the response surface method [1], [42], and involves fitting a low order polynomial via some form of least squares regression. A closely related approach, deriving from geology, is Kriging, whereby Gaussian process models are parameterised by maximum likelihood estimation. A particular example of this is known as the Design and Analysis of Computer Experiments (DACE) model [52], which forms the basis of the EGO algorithm, described in the next section. Optimization methods using response surfaces and/or Kriging meta-models have been successfully used in aeronautical design/engineering applications where only tens or hundreds of function evaluations are possible [57]. In the case of EGO, four low-dimensional multimodal test functions have been optimised to within 1% of optimal in the order of 100 function evaluations [35].

In the evolutionary algorithm community, response surface and Kriging methods have also been used within EAs, to

model either the global landscape [4], [21], [29], [41], [50] or a local region of it [51], in order to reduce the number of expensive function evaluations that need to be carried out during optimization. In addition, similar goals have been obtained using artificial neural networks (either MLPs or RBF networks). Detailed reviews of all these meta-model based EAs can now be found in recent review papers [48], [34], [33], along with new approaches that try to manage the meta-models optimally with respect to the conflicting concerns of optimization and the design of experiments (i.e. improving the model) [47], [48]. Finally, a related but somewhat different approach was proposed recently in [11], in which optimization of an expensive function is tackled by modeling fitness transitions using a finite state machine and, offline, running a number of different evolutionary algorithms on the model, in order to choose the most efficient one. This approach may be attractive when one has many optimization tasks to do on the same or very similar landscapes but a certain investment is needed to build up a sufficiently good model.

Returning to the case of multiple objective functions, we see that a few recent papers have begun to investigate the use of meta-models in MOEAs. Most of these have considered the use of neural networks. The study in [43] demonstrates the use of a neural network approximation combined with the NSGA-II algorithm; the neural net having k outputs, one for each objective function. In this approach, generations using the real evaluation function are alternated with generations using a multi-layer perceptron model of the function derived from the samples collected from the earlier generations. Some speedup is observed over using the original exact fitness function alone, but the study is limited to a single, curve-fitting problem. More recently, [27] and [28] propose and compare two more methods employing neural networks. The first method follows Nain and Deb’s approach, though with a different MOEA being used. The second method is a significant departure from most other meta-modeling schemes, however: an inverse neural network is used to map back from a desired point in the objective space (beyond the current Pareto front) to an estimate of the decision parameters that would achieve it. Test function results presented in [28] for the latter look particularly promising, though fairly long runs (of 20000 evaluations) are considered and it is not clear that on much shorter runs the method would offer significant gains. Finally, a different type of neural network, a self-organizing map, has been employed in [3] to replace standard variation operators with adaptive ones. However, so far, tests have not explicitly compared the performance of this approach with other methods.

Most recently, three papers [20], [22], [23] have appeared that use a Kriging approach based on the DACE model [52] within an MOEA. This approach is the most similar to the one developed independently by the author and described herein.

III. THE EGO ALGORITHM

The EGO algorithm for global optimization of expensive black-box functions was first introduced and described in [35]. It makes use of ‘Kriging’ to model the search landscape from solutions visited during the search. More specifically,

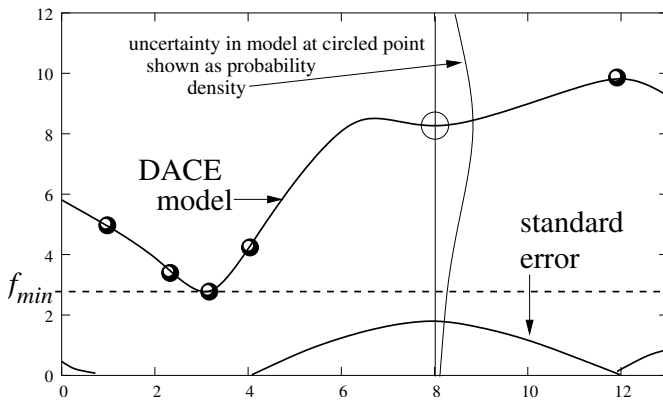


Fig. 1. The uncertainty about the function's value at a point (such as $x=8$ above) can be treated as if the value there were a realization of a normal random variable with mean and standard deviation given by the DACE model and its standard error [Adapted from [35]]

it exploits a version of the *design and analysis of computer experiments* (DACE) model of [52], based on Gaussian processes. We do not give a mathematical description of EGO or DACE here, but simply provide some motivation for its choice in this study, and a brief description of the main procedures involved. The reader is referred to [35] for a fuller explanation.

The DACE model used in Jones [35], like any other method of supervised learning, is subject to certain no-free-lunch theorems [59], similar to those for optimization, that means its general application cannot be recommended on any theoretical basis. Nonetheless, on functions where we expect a degree of local smoothness, where noise is low, and the number of dimensions is not excessive, the DACE model seems to be a good choice for building the kind of approximation that is needed by a search algorithm restricted to a low number of function evaluations. In particular:

- The model always interpolates the points (independently of the chosen parameters of the model)
- The model has a small, fixed number of parameters: $2d+2$ of them, where d is the dimension of the function to be optimized
- The likelihood of the model given the data has a simple closed form expression from which it is possible to compute the maximum likelihood model — giving the approach a sound statistical interpretation
- The error in the expected cost of a solution also has a simple, closed form expression. Thus, the model estimates its own uncertainty.

The last point above is of particular relevance. Knowledge of the error in the response surface is a useful property when searching a cost landscape, and EGO makes use of this property explicitly, as we explain below.

The EGO algorithm begins by first generating a number of solutions in a latin hypercube (i.e., a space-filling design), and by then finding the maximum likelihood DACE model that best explains these solutions (making use of some suitable optimization algorithm). To generate a new solution to evaluate, EGO searches for the solution that maximises what Jones et al [35] call “the expected improvement”—the part of the curve

of the standard error in the model that lies below the best cost sampled so far (see Figure 1). This effectively means that EGO weighs up both the predicted value of solutions, *and the error in this prediction*, in order to find the one that has the greatest potential to improve the minimum cost. EGO does NOT just choose the solution that the model predicts would minimize the cost. Rather, it *automatically* balances exploitation and exploration: where a solution has low predicted cost and low error, it may not be as desirable as a solution whose predicted cost is higher but whose associated error of prediction is also higher. Importantly, the EGO algorithm uses a closed form expression for the expected improvement and it is thus possible to search the decision space globally for the solution that maximises this. Once a new solution has been chosen and evaluated (using the true, expensive cost function), the DACE model is updated with this new information, and the next solution is chosen using this updated model.

Computing the maximum likelihood model at each step of EGO requires several hundred or thousand matrix inversions on a square matrix of dimension equal to the number of solutions so far sampled. As such, it is rather expensive in terms of computation per evaluation, and this overhead increases with time. However, if EGO is efficient with respect to the number of function evaluations needed to reach a given cost, the computational cost can be considered irrelevant in many optimization contexts.

Before we go on to explain how we extend EGO to the multiobjective optimization case, we first give details of our own implementation of the basic algorithm, and some illustrative test results.

A. Implementation of EGO

EGO was implemented in C using the open source matpack library⁴ for coding the matrix operations. The initial solutions are generated using a latin hypercube routine, following a description in Numerical Recipes [49]. The number of initial functions is set to $11d - 1$, where d is the dimension of the function to be optimized, as suggested in [35].

In order to optimise the likelihood function, the Nelder and Mead downhill simplex routine from Numerical Recipes [49] has been used. This is restarted twenty times in order to give a more robust search of the model space. In order to find the best solution to visit next, a genetic algorithm is embedded within EGO, to search the decision parameter space globally. We have not used branch and bound to find the best solution to visit, as suggested by Jones et al., simply because of the difficulty of implementing it, particularly for the more complicated case of constrained optimization, which we would like to tackle in future.

Because the computational overhead increases with each function evaluation, we have found it necessary to cap the size of the correlation matrix (on which the model is based) to have a maximum dimension of 80⁵. This means that beyond the

⁴www.matpack.de

⁵A different library for matrix operations (from <http://cpplapack.sourceforge.net/doc/html/>), based on ATLAS, is being investigated with the hope that it will reduce this problem.

80th function evaluation, 80 solutions are selected at random without replacement from the list of all previously visited solutions — and the DACE model for that iteration is based on these. Potentially, this reduces the accuracy and reliability of EGO (potentially causing it to revisit points). However, beyond matrix sizes of about 80-100, EGO was taking several minutes per evaluation (and this begins to increase dramatically). For testing, this is prohibitive, although in practice a function evaluation can easily take much longer than this. In ParEGO, we also cap the size of the model but use a slightly more advanced selection procedure, as described in Section IV.

B. Testing the implementation of EGO

Testing of our implementation of EGO was carried out on multimodal (epistatic) functions, up to a dimension of about 10, taken from [41]. On each function, we ran EGO (51 times) for just $20d$ evaluations (where d is the dimension of the function). In order to get some idea of the quality of these results, we compare with a random search algorithm (Tables 1 and 2), and with the published results of Liang et al [41], who implement an evolutionary algorithm also with landscape approximation (Table 3). The former results are generally more useful: these allow one to estimate the quality of solutions reached by EGO in terms of the number of solutions in the search space that have that cost or a lower one. The results show that EGO is up to 250 times faster than random search (on four of the functions), more than 50 times faster on a further four, and, at worst, about the same as random search on one highly multimodal function.

The results presented in Table 3 are useful only in that they suggest EGO is performing very well on *some* functions (far better than the EA of Liang et al); but in general it is hard to make comparisons because Liang et al's algorithm is run for 1 to 3 orders of magnitude longer (in terms of evaluations) than we run EGO for.

IV. EXTENDING EGO TO THE MULTIOBJECTIVE CASE: PAREGO

The EGO algorithm could be extended for use with multiobjective optimization problems in a number of different ways. The simplest approach, which we investigate here, converts the k different cost values of a solution into a single cost via a parameterised scalarizing weight vector. By choosing a different (parameterisation of the) weight vector at each iteration of the search, an approximation to the whole Pareto front can be built up gradually.

ParEGO begins by normalizing the k cost functions with respect to the known (or estimated) limits of the cost space, so that each cost function lies in the range $[0,1]$. Then, at each iteration of the algorithm, a weight vector λ is drawn uniformly at random from the set of evenly distributed vectors defined by:

$$\Lambda = \left\{ \lambda = (\lambda_1, \lambda_2, \dots, \lambda_k) \mid \sum_{j=1}^k \lambda_j = 1 \wedge \forall j, \lambda_j = l/s, l \in \{0, \dots, s\} \right\}, \quad (1)$$

with $|\Lambda| = \binom{s+k-1}{k-1}$, so that the choice of s determines how many vectors there are in total. The scalar cost of a solution is then computed using the augmented Tchebycheff function:

$$f_\lambda(x) = \max_{j=1}^k (\lambda_j \cdot f_j(x)) + \rho \sum_{j=1}^k \lambda_j \cdot f_j(x) \quad (2)$$

where ρ is a small positive value which we set to 0.05. It would be possible to use other scalarizing functions — the augmented Tchebycheff function has been chosen here because the non-linear part of the function means that points on non-convex regions of the Pareto front can be minimizers of this function, and thus *non-supported* solutions can be obtained, while the linear part of the function ensures that solutions that are weakly dominated by Pareto optimal solutions are rewarded less than Pareto optimal ones. Using (2) the scalar costs of all previously visited solutions is computed and, using all or a selection of these, a DACE model of the landscape is constructed by maximum likelihood. The solution that maximises the expected improvement with respect to this DACE model is determined. This becomes the next point, and is evaluated on the real, expensive cost function, completing one iteration of ParEGO.

Pseudocode for the entire ParEGO algorithm is given in Algorithm 1. As in our implementation of EGO, the Nelder and Mead downhill simplex algorithm is used (with 20 restarts) to maximize the likelihood of the DACE model (line 25 of Algorithm 1). The evolutionary algorithm used within ParEGO to search for the solution which maximises the expected improvement (line 28) is implemented as follows.

- **Population size:** 20 solutions
- **Population update:** steady state (one offspring produced per generation, from either a crossover or cloning event, followed by a mutation)
- **Generations/evaluations:** 10,000 evaluations
- **Reproductive selection:** binary tournament without replacement
- **Crossover:** simulated binary crossover [18] with probability 0.2, producing one offspring
- **Mutation:** parameter value shifted by $\pm 1/100 \cdot \mu \cdot (\text{parameter range})$, where μ is drawn uniformly at random from $(0.0001, 1)$, and p_m , the per-gene mutation probability, is $1/d$.
- **Replacement:** offspring replaces (first) parent if it is better, else it is discarded
- **Initialization:** 5 solutions are mutants⁶ of the 5 best solutions evaluated on the real fitness function under the prevailing λ vector; the remaining 15 solutions are generated in a latin hypercube in parameter space

In practice, on a very expensive cost function, all solutions evaluated should be used to update the DACE model, at every iteration. However, in our experiments (because of the need to do 21 runs on a large number of functions to collect performance data), we used a simple, heuristic method of choosing a subset of the solutions evaluated to update the

⁶The mutation is carried out as described above except that mutants are checked to ensure they are different from their parents.

TABLE I

RESULTS OF RUNNING OUR IMPLEMENTATION OF EGO ON OUR 13 TEST FUNCTIONS (6 ARE SHOWN ABOVE — THE REMAINING 7 ARE IN TABLE 2).

COLUMN 1 GIVES THE NAME OF THE FUNCTION, WHILE COLUMNS 2 AND 3 GIVE, RESPECTIVELY, ITS DIMENSION d AND ITS NUMBER OF LOCAL OPTIMA. COLUMN 4 SHOWS THE MEDIAN BEST EVALUATION REACHED FROM 51 RUNS OF EGO, WHERE THE NUMBER OF FUNCTION EVALUATIONS USED IN EACH RUN, IS JUST $20d$, WHERE d IS THE DIMENSION OF THE FUNCTION. COLUMNS 5 AND 6 SHOW THE MEDIAN EVALUATION REACHED BY 51 RUNS OF RANDOM SEARCH FOR THE NUMBER OF FUNCTION EVALUATIONS INDICATED. THE MEDIAN VALUE WHICH IS CLOSEST TO EGO'S IS SHOWN IN

BOLD FONT. COLUMNS 7 AND 8 SUMMARISE THE RESULTS OF MANN-WHITNEY RANK-SUM TESTS BETWEEN EGO AND THE RANDOM SEARCH ALGORITHM. THE "EGO WINS" COLUMN GIVES THE NUMBER OF ITERATIONS THAT A RANDOM SEARCH CAN BE RUN (IN TERMS OF A MULTIPLE OF EGO'S NUMBER OF ITERATIONS), AND STILL BE STATISTICALLY BEATEN BY EGO (WITH HIGHEST ASSOCIATED CONFIDENCE LEVEL). THE FINAL COLUMN INDICATES THE SAME THING, BUT WHERE NO DIFFERENCE BETWEEN THE ALGORITHMS' PERFORMANCE CAN BE FOUND (AND THE *lowest* CONFIDENCE LEVEL AT WHICH THE TEST *would* HAVE FOUND A DIFFERENCE HAD THERE BEEN ONE). ON FIVE TEST FUNCTIONS, EGO IS 250 TIMES FASTER THAN RANDOM SEARCH, ON FOUR OTHERS IT IS BETWEEN 50 AND 250 TIMES FASTER. AN INTERPRETATION OF THE RESULTS IS THAT, E.G. ON f_{17} , EGO FINDS A FUNCTION VALUE THAT IS SHARED OR BETTERED BY LESS THAN APPROXIMATELY $1/50000d = 0.001\%$ OF THE ENTIRE SEARCH SPACE, IN JUST 40 FUNCTION EVALUATIONS. EGO GENERALLY PERFORMS WORST ON THE FUNCTIONS WITH VERY MANY LOCAL OPTIMA, THOUGH IT DOES WELL ON BOTH THE GENERALISED GRIEWANK AND SCHWEFEL FUNCTIONS (HERE WITH 6 AND 10 DIMENSIONS, RESPECTIVELY, RATHER THAN

THE USUAL 30 — SEE TABLE 2)

Function	d	$\#opt$	EGO ($20d$ evals)	Random search (#evals)	EGO wins	No diff
f_{14}	2	25	10.7634	54.696730 (20d) 17.379662 (40d) 10.785819 (100d) 2.054400 (1000d) 1.022432 (5000d)	2X (99%)	5X (90%)
f_{15}	4	∞	0.0313316	0.068178 (20d) 0.046495 (40d) 0.025082 (100d) 0.004304 (1000d) 0.002271 (5000d)	2X (99%)	5X (90%)
f_{16}	2	6	-1.03149	-0.003015 (20d) -0.247832 (40d) -0.680195 (100d) -1.005488 (1000d) -1.024494 (5000d)	250X (99%)	—
f_{17}	2	3	0.400733	2.717814 (20d) 1.824595 (40d) 1.522864 (100d) 1.411058 (1000d) 1.401838 (5000d)	100X (99%)	250X (90%)
f_{18}	2	4	4.56758	36.718825 (20d) 22.858280 (40d) 6.779331 (100d) 3.341144 (1000d) 3.112699 (5000d)	5X (95%)	50X (99%)
f_{19}	3	4	-3.86258	-3.556176 (20d) -3.730429 (40d) -3.763582 (100d) -3.835820 (1000d) -3.856405 (5000d)	250X (99%)	—

model, as follows: At each iteration: (i) if the iteration number $iter$ is less than 25, all $11d-1+iter$ solutions evaluated so far, are used to update the model; and (ii) if $iter \geq 25$ a subset of $11d-1+25$ solutions is used, where the first half of them are the best j solutions under the prevailing scalarizing vector λ and the other half are selected at random without replacement. Further details of the parameter settings used in ParEGO are given in Table 5.

V. TEST FUNCTION SUITE

A. Notes on the selection of functions

A number of good attempts at designing test function suites and/or general schemes for test function generation have

been proposed in the multiobjective optimization literature, of which those described in [19], [46], [56] are some of the best. The DTLZ functions in [19] are very popular, partly because they are scalable in the number of objectives and parameters, and partly because each function is accompanied by a description explaining exactly what difficulty it poses to an MOEA, and also allowing this to be easily tuned. However, while it is tempting to employ this suite wholesale, there is (at least) one drawback to it, which leads us to include functions from [56] and [46] as well.

The drawback in Deb et al's suite is that most of the functions work on a scheme where only $k-1$ parameters control the position *parallel* to the Pareto front in the objective

TABLE II

THE RESULTS ON THE REMAINING SIX TEST FUNCTIONS. SEE TABLE I CAPTION. THE NUMBER OF LOCAL OPTIMA IN THE GRIEWANK FUNCTION IS NOT KNOWN — THOUGH FOR A 30-DIMENSIONAL VERSION OF THE SAME FUNCTION IT IS $\simeq 10^{54}$.

Function	d	$\#opt$	EGO (20d evals)	Random search (#evals)	EGO wins	No diff
f_{20}	6	4	-3.30147	-2.016387 (20d) -2.372988 (40d) -2.607558 (100d) -3.011547 (1000d) -3.093685 (5000d)	250X (99%)	—
f_{21}	4	5	-2.43167	-0.536953 (20d) -0.656933 (40d) -0.908331 (100d) -2.296372 (1000d) -3.356989 (5000d)	50X (95%)	100X (90%)
f_{22}	4	7	-2.51587	-0.762930 (20d) -1.029108 (40d) -1.468311 (100d) -2.264707 (1000d) -3.589095 (5000d)	50X (95%)	100X (90%)
f_{23}	4	10	-2.69769	-0.966357 (20d) -1.106497 (40d) -1.463071 (100d) -2.613725 (1000d) -3.741701 (5000d)	50X (90%)	100X (90%)
f_{24}	2	∞	0.345967	0.319360 (20d) 0.203386 (40d) 0.179859 (100d) 0.038538 (1000d) 0.014670 (5000d)	—	1X (90%)
Griewank	6	?	1.0205	41.558208 (20d) 30.257815 (40d) 24.441566 (100d) 11.717705 (1000d) 6.968934 (5000d)	250X (99%)	—
Schwefel	10	$> 10^9$	-4186.97	-3745.924 (20d) -3812.0805 (40d) -3900.232 (100d) -4025.030 (1000d) -4079.477 (5000d)	250X (99%)	—

TABLE III

COMPARISON OF EGO WITH PUBLISHED RESULTS OF RUNNING EANA [41] ON THE 11 LOWER-DIMENSIONAL FUNCTIONS. FOR EGO, WE RAN FOR EXACTLY $20d$ FUNCTION EVALUATIONS (COLUMN 2) AND QUOTE THE MEAN AND BEST COST FOUND FROM 51 RUNS (COLUMNS 3 AND 4, RESPECTIVELY). FOR EANA, THE ALGORITHM WAS STOPPED WHEN IT REACHED EITHER THE GLOBAL OPTIMUM (UP TO THE PRECISION OF THE COMPUTER) OR SOME UNSPECIFIED EVALUATION LIMIT [41]. THUS, THE MEAN NUMBER OF EVALUATIONS IS REPORTED (COLUMN 5) AND THE MEAN COST OF THE SOLUTION OVER THE 50 RUNS PERFORMED (LAST COLUMN). COMPARISON IS DIFFICULT BUT ON SOME FUNCTIONS, EGO OCCASIONALLY REACHES COMPARABLE FITNESS LEVELS ONE OR TWO ORDERS OF MAGNITUDE FASTER THAN THE MEAN NUMBER OF FUNCTION EVALUATIONS QUOTED FOR EANA. (EANA IS AN EVOLUTION STRATEGY, ALSO BASED ON BUILDING LANDSCAPE APPROXIMATIONS)

Function	EGO			EANA	
	#evals	mean	best	mean #evals	mean
f_{14}	40	20.7267	0.998029	1685	7.36
f_{15}	80	0.0309024	0.00157531	31645	0.000307
f_{16}	40	-1.03019	-1.031628	863	-1.031628
f_{17}	40	0.438083	0.3979	945	0.3979
f_{18}	40	12.8115	3.00007	822	3
f_{19}	60	-3.8624	-3.86278	1297	-3.86
f_{20}	120	-3.2562	-3.32135	8357	-3.3128
f_{21}	80	-2.72465	-6.60786	2958	-9.6512
f_{22}	80	-2.70663	-7.66332	2880	-9.1430
f_{23}	80	-3.32506	-10.276	3118	-9.5767
f_{24}	40	0.330584	0.0393655	1181	0.0297

Algorithm 1 ParEGO pseudocode

```

1: procedure PAREGO( $f, d, k, s$ )
2:    $xpop[] \leftarrow \text{LATINHYPERCUBE}(d)$  /* Initialize using
   procedure defined in line 15 */
3:   for each  $i$  in 1 to  $11d - 1$  do
4:      $zpop[i] \leftarrow \text{EVALUATE}(xpop[i], f)$  /* (line 36) */
5:   end for
6:   while not finished do
7:      $\lambda \leftarrow \text{NEWLAMBDA}(k, s)$  /* (line 19) */
8:      $model \leftarrow \text{DACE}(xpop[], zpop[], \lambda)$  /* (line 22) */
9:      $xnew \leftarrow \text{EVOLALG}(model, xpop[])$  /* (line 28)
       */
10:     $xpop[] \leftarrow xpop[] \cup \{xnew\}$ 
11:     $znew \leftarrow \text{EVALUATE}(xnew, f)$ 
12:     $zpop[] \leftarrow zpop[] \cup \{znew\}$ 
13:  end while
14: end procedure

15: procedure LATINHYPERCUBE( $d$ )
16:  divide each dimension of search space into  $11d - 1$ 
   'rows' of equal width
17:  return  $11d - 1$  vectors  $x$  such that no two share the
   same 'row' in any dimension
18: end procedure

19: procedure NEWLAMBDA( $k, s$ )
20:  return a  $k$ -dimensional scalarizing weight vector cho-
   sen uniformly at random from amongst all those
   defined by equation 1
21: end procedure

22: procedure DACE( $xpop[], zpop[], \lambda$ )
23:  compute the scalar fitness  $f_\lambda$  of every cost vector in
    $xpop[]$ , using equation 2)
24:  choose a subset of the population based on the com-
   puted scalar fitness values
25:  maximize the likelihood of the DACE model for the
   chosen population subset
26:  return the parameters of the maximum likelihood
   DACE model
27: end procedure

28: procedure EVOLALG( $model, xpop[]$ )
29:  initialise a temporary population of solution vectors,
   some as mutants of  $xpop[]$  and others purely
   randomly
30:  while set number of evaluations not exceeded do
31:    evaluate the expected improvement of solutions
   using the model
32:    select, recombine and mutate to form new popula-
   tion
33:  end while
34:  return best evolved solution
35: end procedure

36: procedure EVALUATE( $x, f$ )
37:  call the expensive evaluation function  $f$  with the
   solution vector  $x$ 
38:  return true cost vector  $z$  of solution  $x$ 
39: end procedure

```

space, with all the remaining parameters controlling distance *to* the front (i.e., the problems are separable). From this, a natural bias is given to recombination-based EAs because, once a few points on the true Pareto front are found, recombination tends to propagate and preserve the parameters which are common to them (which is all of them, bar the first $k - 1$), while searching over the other parameters, leading to a very fast spreading out over the entire front. Unfortunately, many real-world multiobjective optimization problems are not structured like this; rather, each different optimal tradeoff solution requires changing all or many of the decision parameters.

In [46], similar observations in respect of the DTLZ functions have been made, and furthermore, it was also observed that the Pareto fronts of many other test functions consist of piecewise linear curves and/or single points in the parameter space. The paper goes on to propose a scheme for generating functions that, like [19], allows some different problem difficulties to be incorporated and tuned, combined with a greater control of the curves/surfaces in parameter space comprising the Pareto front. Two example functions using this general scheme are given, and it is these that we employ in our suite.

The functions we employ from [56] are not part of an overall tunable scheme but have been popular in the literature. VLMOP2 has a concave Pareto front while VLMOP3's Pareto front is both nonlinear and asymmetric.

Overall, the final selection of nine test functions includes functions from two to eight decision parameters; functions with a very low density of solutions at the Pareto front; functions with locally optimal Pareto fronts; functions where the Pareto set follows a complicated curve in the parameter space; functions where the Pareto front is disconnected in objective space; and functions where the density of points parallel to the Pareto front is non-uniformly distributed. There is thus a good deal of variety in the difficulties that they pose. We have nonetheless been restrictive in some particular aspects: all functions are unconstrained and while difficult, are not overly high-dimensional (in parameter space), and have a reasonable, rather than pathological degree of ruggedness. And, we have kept to functions of two and three objectives only. These restrictions accord with our description (in Section 1) of certain kinds of expensive engineering/scientific problem, where we hope to obtain good results in a very small number of function evaluations.

In the following we describe each of our test functions in turn.

KNO1

This test function is the only one not borrowed from the existing literature. There are just two parameters and two objectives in the problem:

$$\text{Minimize } f_1 = 20 - r \cos(\phi),$$

$$\text{Minimize } f_2 = 20 - r \sin(\phi),$$

$$r = 9 - [3 \sin(5/2(x_1 + x_2)^2) + 3 \sin(4(x_1 + x_2)) + 5 \sin(2(x_1 + x_2) + 2)],$$

$$\phi = \pi/12(x_1 - x_2 + 3),$$

$$x_1, x_2 \in [0, 3].$$

The distance from the Pareto front, controlled by r , is a function of the sum of the two decision parameters, while the location transverse to the Pareto front is controlled by the difference between the two parameters. Thus, each point on the Pareto front is unique in both parameters: the Pareto set is made up of all pairs of parameters that sum to 4.4116. The true Pareto front has a greater extent than, and lies just beyond, a locally optimal Pareto front with a much larger basin of attraction. Altogether, there are fifteen locally optimal fronts in this function.

OKA1

The OKA1 test function [46] is defined as:

$$\begin{aligned} \text{Minimize } f_1 &= x'_1, \\ \text{Minimize } f_2 &= \sqrt{2\pi} - \sqrt{|x'_1|} + 2|x'_2 - 3 \cos(x'_1) - 3|^{\frac{1}{3}}, \\ x'_1 &= \cos(\pi/12)x_1 - \sin(\pi/12)x_2, \\ x'_2 &= \sin(\pi/12)x_1 + \cos(\pi/12)x_2, \\ x_1 &\in [6 \sin(\pi/12), 6 \sin(\pi/12) + 2\pi \cos(\pi/12)], \\ x_2 &\in [-2\pi \sin(\pi/12), 6 \cos(\pi/12)]. \end{aligned}$$

The Pareto optima (in the parameter space) lie on the curve, $x'_2 = 3 \cos(x'_1) + 3$ ($x'_1 \in [0, 2\pi]$). In addition the density of solutions falls away near to the Pareto front.

OKA2

The OKA2 test function [46] is defined as:

$$\begin{aligned} \text{Minimize } f_1 &= x_1, \\ \text{Minimize } f_2 &= 1 - \frac{1}{4\pi^2}(x_1 + \pi)^2 + |x_2 - 5 \cos(x_1)|^{\frac{1}{3}} + \\ &|x_3 - 5 \sin(x_1)|^{\frac{1}{3}} \\ x_1 &\in [-\pi, \pi], \quad x_2, x_3 \in [-5, 5]. \end{aligned}$$

The Pareto optima lie on a spiral-shaped curve in the 3-dimensional parameter space. In addition the density of solutions falls away steeply near to the Pareto front.

VLMOP2

Veldhuizen and Lamont's MOP2 function [56] is scalable in the number of decision variables. We use just two:

$$\begin{aligned} \text{Minimize } f_1 &= 1 - \exp\left(-\sum_{\{i=1, \dots, n\}} (x_i - 1/\sqrt{n})^2\right), \\ \text{Minimize } f_2 &= 1 + \exp\left(-\sum_{\{i=1, \dots, n\}} (x_i - 1/\sqrt{n})^2\right), \\ x_1, x_2 &\in [-2, 2], n = 2. \end{aligned}$$

The Pareto front is concave and the Pareto optima lie on the diagonal passing from $-1/\sqrt{n}, -1/\sqrt{n}$ to $1/\sqrt{n}, 1/\sqrt{n}$ in parameter space.

VLMOP3

Veldhuizen and Lamont's MOP3 [56] is a three-objective function of two decision variables:

$$\begin{aligned} \text{Minimize } f_1 &= 0.5(x^2 + y^2) + \sin(x^2 + y^2), \\ \text{Minimize } f_2 &= \frac{(3x - 2y + 4)^2}{8} + \frac{(x - y + 1)^2}{27} + 15, \\ \text{Minimize } f_3 &= \frac{1}{x^2 + y^2 + 1} - 1.1 \exp(-x^2 - y^2), \\ x, y &\in [-3, 3]. \end{aligned}$$

This test function has a disconnected Pareto optimal set, and the Pareto front is a curve 'following a convoluted path through objective space'.

DTLZ1a

We have selected four of the test functions described in [19], and have made alterations to them in order to keep the number of decision variables small, and/or to reduce the ruggedness to 'reasonable' levels. We append an "a" to the name of these functions to show that we have changed them slightly.

The first function from [19], DTLZ1, is used with two objectives and six decision variables here. This gives the following function, DTLZ1a:

$$\begin{aligned} \text{Minimize } f_1 &= \frac{1}{2}x_1(1 + g), \\ \text{Minimize } f_2 &= \frac{1}{2}(1 - x_1)(1 + g), \\ g &= 100 \left[5 + \sum_{i \in \{2, \dots, 6\}} (x_i - 0.5)^2 - \cos(2\pi(x_i - 0.5)) \right], \\ x_i &\in [0, 1], i \in \{1, \dots, n\}, n = 6. \end{aligned}$$

The ruggedness of the function is controlled by the g function. In the original DTLZ1, 20π is used in the cosine term, giving a total of $11^k - 1$ local Pareto optimal fronts. This is excessively rugged for our purposes, hence we use 2π instead. The Pareto optimal set consists of all solutions where all but the first decision variables are equal to 0.5, and the first decision variable may take any value in $[0, 1]$.

DTLZ2a and DTLZ4a

The next two test problems have three objective functions defined on eight decision variables. DTLZ2a and DTLZ4a are given by:

$$\begin{aligned} \text{Minimize } f_1 &= (1 + g) \cos(x_1^\alpha \pi/2) \cos(x_2^\alpha \pi/2), \\ \text{Minimize } f_2 &= (1 + g) \cos(x_1^\alpha \pi/2) \sin(x_2^\alpha \pi/2), \\ \text{Minimize } f_3 &= (1 + g) \sin(x_1^\alpha \pi/2), \\ g &= \sum_{i \in \{3, \dots, 8\}} (x_i - 0.5)^2, \\ x_i &\in [0, 1], i \in \{1, \dots, n\}, n = 8, \\ \alpha &= \begin{cases} 1, & \text{if function is DTLZ2a} \\ 100, & \text{if function is DTLZ4a.} \end{cases} \end{aligned}$$

The Pareto front is one eighth of a sphere of radius 1, centred on $0, 0, 0$. The Pareto optimal set consists of all solutions where all but the first decision variables are equal to 0.5, and the first decision variable may take any value in $[0, 1]$. The effect of setting $\alpha = 100$ is to severely bias the density distribution of solutions toward the $f_3 - f_1$ and $f_2 - f_1$ planes.

DTLZ7a

This problem has four disconnected regions in the Pareto front (in objective space):

$$\text{Minimize } f_1 = x_1,$$

$$\text{Minimize } f_2 = x_2,$$

$$\text{Minimize } f_3 = (1 + g)h,$$

$$g = 1 + 9/6 \sum_{i \in \{3, \dots, 8\}} x_i,$$

$$h = 3 - \sum_{i \in \{1, 2\}} \left[\frac{f_i}{1 + g} (1 + \sin(3\pi f_i)) \right],$$

$$x_i \in [0, 1], i \in \{1, \dots, n\}, n = 8.$$

VI. SELECTED PERFORMANCE ANALYSIS TECHNIQUES

The output of a multiobjective optimizer for a single run is an approximation set: the set of all nondominated points found during the run. Assessing the performance of optimizers thus relies upon some means of measuring or comparing the quality of one or more approximation sets. Many such *performance indicators*⁷ exist in the literature but it has been shown in some recent works [38], [39], [61] that several of these are not consistent with the partial ordering of approximation sets that follows logically from the basic principles of Pareto optimization. For example, in the worst case, two sets of nondominated points, A and B may be compared, where *every* point in B is dominated by at least one point in A , and yet B yields a better evaluation under the chosen indicator [39]. This is particularly the case for indicators that try to assess one isolated aspect, such as the diversity, extent or proximity to the Pareto front, of a nondominated set. Fortunately, indicators that simultaneously take into account *all* of these aspects of what it means to be a good nondominated set do exist, and moreover they are compatible and/or complete with respect to the true partial ordering of nondominated sets. Two such ‘well-behaved indicators’, which we choose to use in our experiments, are the hypervolume indicator and the binary epsilon indicator, described below. For a more detailed discussion of assessment methods, the reader is referred to [61], where a full classification of indicators is given.

In addition to choosing appropriate performance indicators, one must also decide how to collate and interpret the results from multiple runs, and how to present results graphically, if at all. Since, in our optimization scenario, it is likely that there would be just one (or, perhaps, two) shots at performing the optimization, the average-case and worst-case performance

⁷We prefer the term ‘indicator’, following [61], because the terms ‘measure’ and ‘metric’ have reserved meanings in mathematics that not all indicators conform to.




Fig. 2. Five attainment surfaces are shown, representing the output of five runs of an optimizer. The two diagonal lines intersect the five surfaces at various points; in both cases, the circle indicates the intersection that weakly dominates at least $5 - 3 + 1 = 3$ surfaces and is also weakly dominated by at least 3 surfaces. Therefore, these two points both lie on the third summary attainment surface

are paramount. Hence, for each, test function the following information is given:

- the mean/median results of individual runs (as well as standard deviation/interquartile ranges)
- the statistical significance of differences in the distributions
- for the 2-objective functions only, plots of the median, and worst attainment surfaces achieved
- for the 3-objective functions only, plots of the worst attainment surfaces achieved only (since 3d plots showing other surfaces are visually confusing).

We give all details in the following sections and also make our assessment methods source code available for public use at [37].

A. Hypervolume indicator (a.k.a the \mathcal{S} measure) [60]

This indicator assesses the size (hypervolume or Lebesgue integral) of the region dominated by a sample set of points, thus larger values generally indicate better nondominated sets. The region itself must be bounded from above in some way, and for this some point b is chosen, which must itself be dominated by every point in the sample set.

The hypervolume indicator \mathcal{S} is well-behaved in the sense that given two sets, A and B , and using the same bounding point, if $\mathcal{S}(A) > \mathcal{S}(B)$ then it cannot be the case that B is better than A (where better is used in the sense defined in [61]). However, it must be noted that this result does not imply that A is better than B , either. Nonetheless, in the absence of preferences, and with a reasonable choice of bounding point, the \mathcal{S} measure gives results that concur with intuition — proximity to the true PF, extent and evenness of solutions all tend to be rewarded.

In order to choose a bounding point for application of the \mathcal{S} measure, we use the following method. First, the collection of nondominated point sets from all runs of all algorithms are




Fig. 3. Five sets of nondominated points and the best, median and worst attainment surfaces that they define. The interpretation of the median attainment surface is that, for every point on it (independently), a point (weakly) dominating this was obtained in at least 50% of the nondominated sets. Similarly, the worst attainment surface indicates the goals achieved in 100% of the sets. The best attainment surface indicates the goals achieved by the aggregation of all sets. In this study, the best attainment surface is irrelevant and is not included in plotted results

aggregated into a single superset. Then, the ideal and the anti-ideal point of this superset is found. The bounding point is then the anti-ideal point shifted by δ times the range, in each objective:

$$\mathbf{b} = (b_1, b_2, \dots, b_k), \quad \text{with} \\ b_j = \max_j + \delta(\max_j - \min_j), \quad j \in \{1, \dots, k\},$$

where \max_j and \min_j are the maximum and minimum value, respectively, on the j th objective, found within the superset. We use $\delta = 0.01$ here.

For the analysis of multiple runs, we compute the \mathcal{S} measure of each individual run, and report the mean and the standard deviation of these. Since the distribution of ParEGO and NSGA-II's results are not necessarily normal, we use the Mann-Whitney rank-sum test to indicate if there is a statistically significant difference in the position of the two distributions.

B. Additive binary epsilon indicator

There are two ways in which the hypervolume indicator, used on its own, *can* be misleading. First, there is no way, from looking at \mathcal{S} values in isolation, of inferring whether one set is actually better than another in a strict sense. Secondly, the choice of bounding point is rather arbitrary, and this can affect the ordering of some pairs of sets.

A method that avoids these particular difficulties is the additive binary epsilon indicator [61]. This takes a pair of nondominated sets A and B and returns a pair of numbers,

$$(I_A, I_B)$$

$$I_A = I_{\epsilon+}(A, B) = \inf_{\epsilon \in \mathbb{R}} \{ \forall z^2 \in B \exists z^1 \in A : z^1 \preceq_{\epsilon+} z^2 \} \\ I_B = I_{\epsilon+}(B, A) = \inf_{\epsilon \in \mathbb{R}} \{ \forall z^2 \in A \exists z^1 \in B : z^1 \preceq_{\epsilon+} z^2 \}$$

where $z^1 \preceq_{\epsilon+} z^2$ iff $\forall j$ in $\{1, \dots, k\} : z^1_j \leq \epsilon + z^2_j$, assuming minimization. (Note, $z^1 \preceq_{\epsilon+} z^2$ is read as z^1 epsilon-dominates z^2). A pair of numbers ($I_A \leq 0, I_B > 0$) indicates that A is strictly better than B [61]. A pair of numbers ($I_A > 0, I_B > 0$) indicates that neither set is strictly better than the other — they are incomparable [61]. However, if I_A is less than I_B , then in a weaker sense, it is better because the minimum ϵ value needed so that approximation set A ϵ -dominates B is smaller than the ϵ value needed for B to ϵ -dominate A . In the results we present, unnormalized objective values have been used in the computation of I_A and I_B .

To summarise the results of multiple runs, we compute the binary epsilon indicator for each pair of runs in turn and then report on the median and interquartile range values of I_A and I_B . We prefer the median and IQR to the mean and SD because if the median for either algorithm is less than zero, then this has a precise interpretation — that it was strictly better than the other algorithm on at least 50% of runs. We again use the Mann-Whitney rank sum test to decide if the distribution of I_A values is different from the distribution of I_B values.

C. Median and worst attainment surface plots

An *attainment surface* is the surface uniquely determined by a set of nondominated points that divides the objective space into the region dominated by the set and the region that is not dominated by it [26]. Given n runs of an algorithm, it would be nice to summarise the n results attainment surfaces obtained, using only one or two *summary* surfaces. Such summary attainment surfaces can be defined by imagining a diagonal line in the direction of increasing objective values cutting through the n results attainment surfaces (see Figure 2). The intersection on this line that weakly dominates at least $n - p + 1$ of the surfaces and is weakly dominated by at least p of them, defines one point on the ' p th summary attainment surface'. This surface is the union of all the goals that have been attained in at least p runs (independently).

Obviously, from the definition of the p th summary attainment surface, it is possible to define a median (summary) surface (for an odd number of approximation sets) and a best and worst (summary) surface. For illustration, we plot five approximation sets and their sample median, best and worst attainment surfaces in Figure 3. Notice that the interpretation of the median surface is that, for every point on it, a point (weakly) dominating this was obtained in more than half of the algorithm runs. Similarly, the worst attainment surface has the interpretation that in every algorithm run the entire surface was (weakly) dominated. Thus, plots of these surfaces provide more information than a mere scatter plot of the aggregation of (nondominated) points found from several runs.

For the two-objective problems in this paper, we give the median and worst attainment surface of ParEGO and of

TABLE IV

NSGA-II PARAMETER SETTINGS, WHERE d IS THE NUMBER OF DECISION
PARAMETER DIMENSIONS

<i>Parameter</i>	<i>setting</i>
Population size	20
Maximum generations	13 [100]
Crossover probability	0.9
Real-value mutation probability	$1/d$
Real-value SBX parameter	10
Real-value mutation parameter	50

TABLE V

ParEGO PARAMETER SETTINGS, WHERE d IS THE NUMBER OF DECISION
PARAMETER DIMENSIONS

<i>Parameter</i>	<i>setting</i>
Initial pop. in latin hypercube	$11d - 1$
Total maximum evaluations	250
Number of scalarizing vectors	11 ($k = 2$), 15 ($k = 3$)
Scalarizing function	augmented Tchebycheff
Internal GA evals per iteration	200 000
Crossover probability	0.2
Real-value mutation probability	$1/d$
Real-value SBX parameter	10
Real-value mutation parameter	50

NSGA-II on the same plot, with ParEGO's surfaces shown in solid and NSGA-II surfaces shown in dashed lines. For the three-objective problems, we show just the worst attainment surface of the two algorithms, each on separate plots but with the same axes. The worst attainment surface gives a very rough indication of the worst-case performance of the two algorithms, although it has no statistical significance.

To actually compute the points plotted, we use an algorithm based on the one described in [53], and available from [37].

VII. EXPERIMENTAL DETAILS

To evaluate NSGA-II and ParEGO on the test suite, each algorithm is run 21 times, and all solutions visited are stored. The nondominated sets achieved after a particular number of function evaluations can then be determined and used to estimate performance.

Performance was also compared with a random search run 21 times for up to 10000 function evaluations on all problems.

A. NSGA-II parameter settings

The implementation of NSGA-II that we use is the one available for download from Deb's KANGAL webpage. We use default settings (Table IV), as supplied, with one exception. The population size is reduced to 20 in order to maximise the performance over the short run experiments. This size was determined from some runs in which sizes from 10 to 50 were tried. With a population size of 20, 13 generations gives 260 evaluations in all. This is comparable to the 250 evaluations given to ParEGO. The performance of the two algorithms is also compared at 100 evaluations.

B. ParEGO parameter settings

The main choices in the parameter settings of ParEGO concern the number of scalarizing vectors to use (specified via the parameter s in Equation 1), and in the parameters of the internal EA. For the former we chose a number which would allow several 'passes' over each scalarizing vector. For the evolutionary algorithm within ParEGO, a small number of preliminary experiments were carried out to determine settings that led to robust optimization of the expected improvement of the DACE model. The full set of parameter settings used in all runs of ParEGO are given in Table V.

VIII. RESULTS

Compared to a random search of the space, NSGA-II does well on most problems, even though only 100 or 260 function evaluations are used. An impression of this can be had from the plots (Figures 4 and 5) showing the median and worst attainment surfaces for NSGA-II, compared with 1000 random solutions, plotted in objective space, on four of the test functions. Plots of the attainment surfaces for the random search algorithm, not shown here, suggest that NSGA-II is better than random search over 260 function evaluations on all the test problems.

Tables VI and VII give the results of applying the S measure to the 21 runs of ParEGO and NSGA-II after, respectively, 100 and 250 (260 for NSGA-II) function evaluations. From these results it is clear that ParEGO consistently dominates a larger region of the objective space (modulo the chosen bound point) at 250 evaluations. Not only is the mean higher, but the standard deviation is consistently lower, sometimes by several orders of magnitude, suggesting a much more reliable, and better worst-case, performance. The Mann-Whitney test shows that the differences between the two populations are significant at the 99% confidence level, on *all* functions.

At 100 function evaluations, the outcome is much the same, except for the functions DTLZ4a and DTLZ7a, where NSGA-II obtains a higher mean value (and is significantly better on DTLZ7a). A possible explanation for the relatively poorer performance of ParEGO on these two functions is that it has actually only 'chosen' 13 of the solutions evaluated — the first 87 solutions being the latin hypercube — since these functions both have 8 decision variables. If one were interested in a strong performance after 100 evaluations on these functions, a sparser latin hypercube may be worth considering.

Tables VIII and IX give the equivalent results for the additive binary epsilon indicator. The results are largely in agreement with those of the S measure, despite the fact that the method of measurement is quite different. Once again it is evident that ParEGO has less variation, and that after 250 (260 for NSGA-II) evaluations it is not significantly behind on any problem. Furthermore, on one problem, OKA2, its median indicator score is negative, indicating that its nondominated sets are better in a strict sense than NSGA-II's on more than 50% of runs. On two or three other problems its median indicator value is also very close to zero, and we can confirm that on these problems its nondominated sets are strictly better than NSGA-II's on several of the runs.

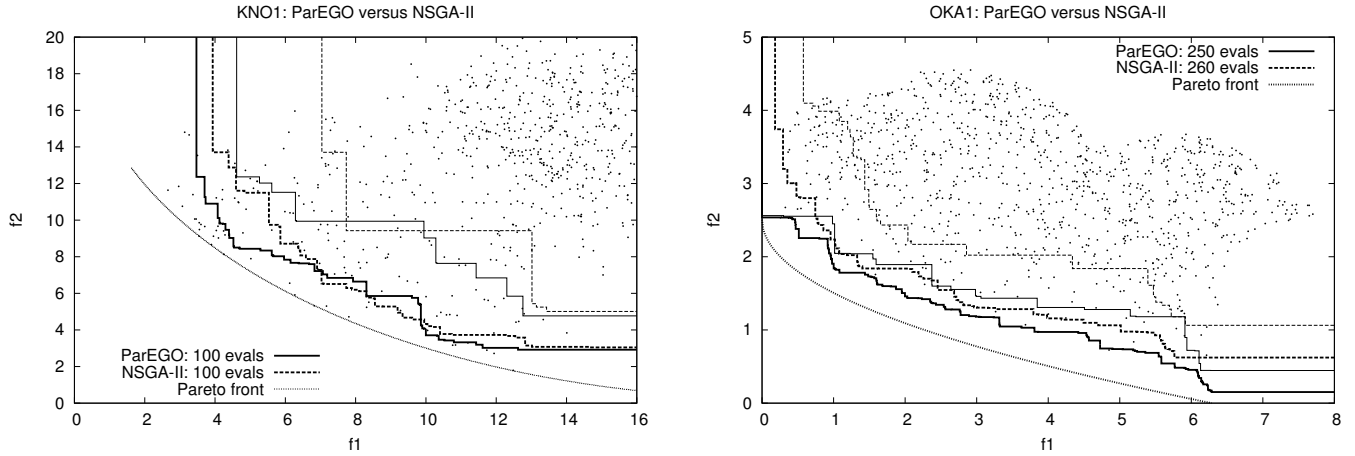


Fig. 4. Attainment surface plot with 1000 random search points also shown, on the KNO1 function after 100 function evaluations (left), and on the OKA1 function after 250/260 function evaluations (right)

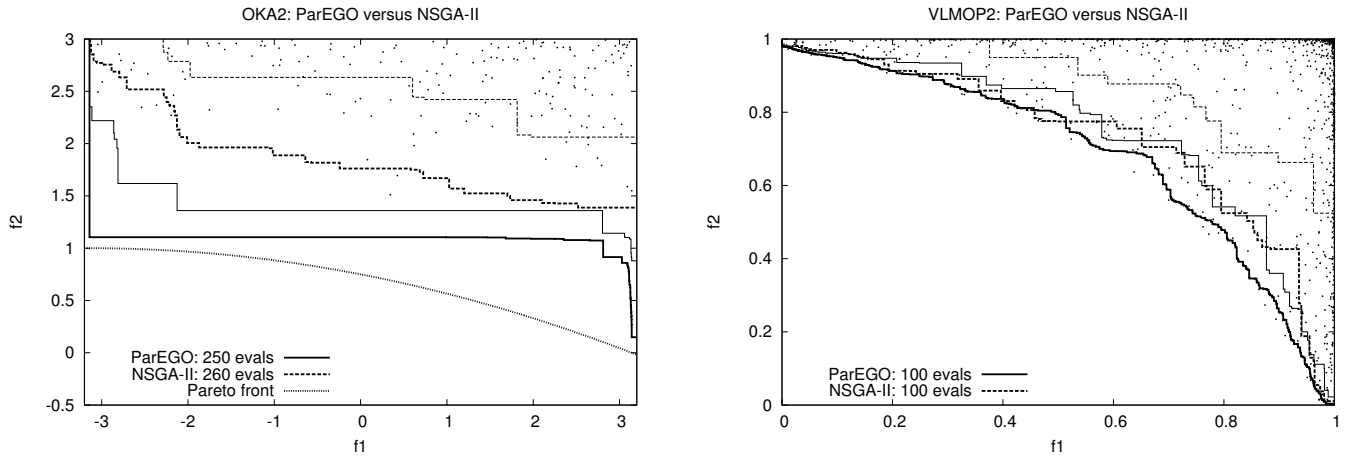


Fig. 5. Attainment surface plot with 1000 random points also shown, on OKA2 after 250/260 function evaluations (left), and on VLMOP2 after 100 function evaluations (right)

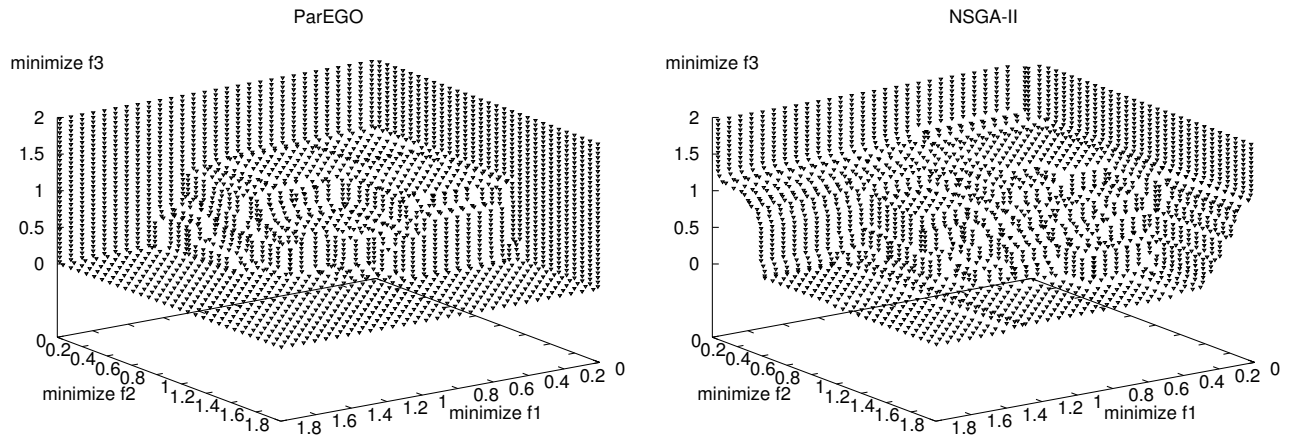


Fig. 6. Worst attainment surfaces after 250/260 function evaluations for ParEGO (left) and NSGA-II (right) on the DTLZ2a function




Fig. 7. Worst attainment surfaces after 250/260 function evaluations for ParEGO (left) and NSGA-II (right) on the DTLZ4a function




Fig. 8. Worst attainment surfaces after 250/260 function evaluations for ParEGO (left) and NSGA-II (right) on the DTLZ7a function

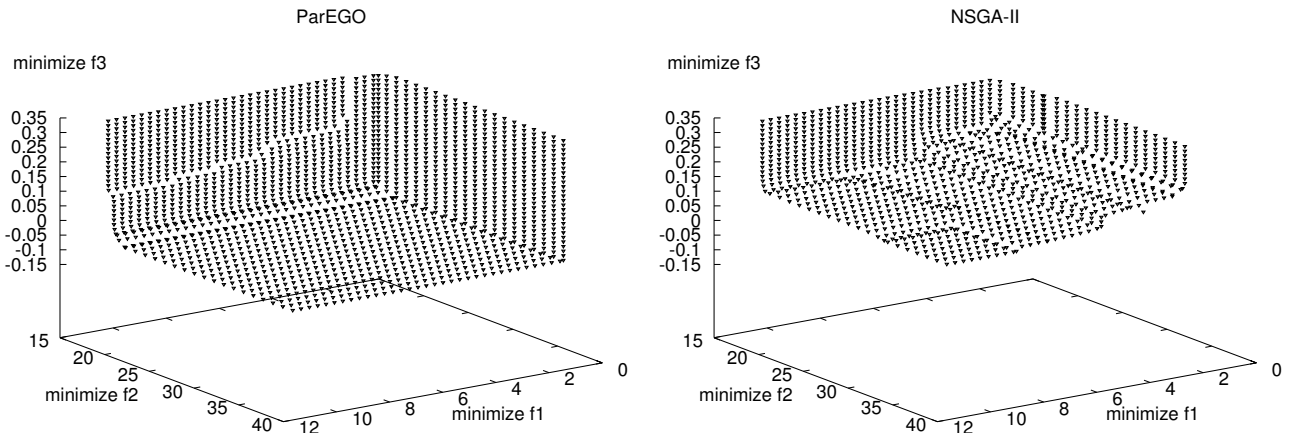


Fig. 9. Worst attainment surfaces after 250/260 function evaluations for ParEGO (left) and NSGA-II (right) on the VLMOP3 function

TABLE VI

VALUES OF THE \mathcal{S} MEASURE AFTER 100/100 EVALUATIONS OF PAREGO/NSGA-II. LARGER VALUES INDICATE BETTER PERFORMANCE. THE DISTRIBUTIONS OF THE \mathcal{S} VALUES ARE TESTED USING THE MANN-WHITNEY RANK SUM TEST. THE z VALUES AND SIGNIFICANCE LEVEL ARE INDICATED. PAREGO IS SIGNIFICANTLY BETTER THAN NSGA-II UNLESS STATED

Function	ParEGO mean (SD)	NSGA-II mean (SD)	z -value	significance
KNO1	94.5363 (7.28313)	87.7112 (8.28212)	-2.427529	> 99%
OKA1	16.9725 (0.355216)	14.138 (1.25082)	-5.546841	> 99%
OKA2	22.2165 (0.887676)	15.5118 (1.42961)	-5.546841	> 99%
VLMOP2	0.320038 (0.00563326)	0.263637 (0.0387093)	-5.446218	> 99%
VLMOP3	46.3881 (0.157458)	34.8668 (12.5845)	-5.320439	> 99%
DTLZ1a	189262 (209.001)	186855 (1943.03)	-5.270128	> 99%
DTLZ2a	4.40784 (0.0906379)	4.06872 (0.172994)	-4.943103	> 99%
DTLZ4a	0.673329 (0.263526)	0.824647 (0.408542)	-0.943340	- no winner
DTLZ7a	12.9895 (0.655795)	16.8967 (1.00974)	-5.546841	> 99% NSGA-II wins

TABLE VII

VALUES OF THE \mathcal{S} MEASURE AFTER 250/260 EVALUATIONS OF PAREGO/NSGA-II. LARGER VALUES INDICATE BETTER PERFORMANCE. THE DISTRIBUTIONS OF THE \mathcal{S} VALUES ARE TESTED USING THE MANN-WHITNEY RANK SUM TEST. THE z VALUES AND SIGNIFICANCE LEVEL ARE INDICATED. PAREGO IS BETTER IN ALL CASES

Function	ParEGO mean (SD)	NSGA-II mean (SD)	z -value	significance
KNO1	108.478 (5.83982)	103.2 (7.92597)	-2.276594	> 99%
OKA1	18.0834 (0.372983)	15.982 (0.681422)	-5.546841	> 99%
OKA2	24.0278 (0.467227)	19.1279 (1.44751)	-5.546841	> 99%
VLMOP2	0.312768 (0.00398254)	0.30501 (0.0119822)	-2.528152	> 99%
VLMOP3	89.9313 (0.113778)	83.0951 (18.8421)	-2.578463	> 99%
DTLZ1a	64869.6 (6.98978)	64682.7 (154.84)	-3.710472	> 99%
DTLZ2a	4.46196 (0.0276682)	4.07554 (0.145971)	-5.546841	> 99%
DTLZ4a	2.14206 (0.478063)	1.42626 (0.809588)	-3.106734	> 99%
DTLZ7a	15.8421 (0.431151)	13.653 (0.930073)	-5.395906	> 99%

TABLE VIII

MEDIAN AND INTERQUARTILE RANGE VALUES OF THE BINARY EPSILON INDICATOR AFTER 100/100 EVALUATIONS OF PAREGO/NSGA-II. LOWER VALUES INDICATE BETTER PERFORMANCE. THE DISTRIBUTIONS OF THE I VALUES ARE TESTED USING THE MANN-WHITNEY RANK SUM TEST. THE z VALUES AND SIGNIFICANCE LEVEL ARE INDICATED. PAREGO IS SIGNIFICANTLY BETTER THAN NSGA-II UNLESS STATED

Function	median values	IQR values	z -value	significance
	$(I_{ParEGO}, I_{NSGA-II})$	$(I_{ParEGO}, I_{NSGA-II})$		
KNO1	(1.5261, 2.2245)	(1.2633, 1.3428)	-2.1508	> 98%
OKA1	(0.3581, 0.9120)	(0.2395, 0.3714)	-5.4965	> 99%
OKA2	(-0.0123, 1.3692)	(0.0540, 0.4530)	-5.5468	> 99%
VLMOP2	(0.0313, 0.1029)	(0.0283, 0.2180)	-5.0186	> 99%
VLMOP3	(0.0781, 0.6956)	(0.2167, 1.0073)	-4.3645	> 99%
DTLZ1a	(0.3168, 17.7269)	(0.8050, 17.2216)	-5.5846	> 99%
DTLZ2a	(0.00780, 0.0613)	(0.0151, 0.0969)	-3.458914	> 99%
DTLZ4a	(0.2276, 0.4153)	(0.4253, 0.7326)	-1.3207	> 90%
DTLZ7a	(0.7813, 0.0963)	(0.6198, 0.9580)	-3.0564	> 99% NSGA-II wins

The plots of the median and worst attainment surfaces (Figures 4–9) tend to confirm the outcome of the measures, above. In particular, the worst attainment surface for NSGA-II is much worse than for ParEGO on several of the problems.

Finally, Figure 10 shows the decision space solutions, and corresponding points in objective space, visited by NSGA-II and ParEGO on the first run on function OKA1. This plot illustrates a considerable difference in the ways that ParEGO and NSGA-II sample the decision space, which has also been

observed in other decision space plots: NSGA-II is much more aggressive at honing in on good points, exploiting them much more than ParEGO. However, despite ParEGO's greater exploration, it still manages to exploit the good solutions it has found very effectively, and so returns both very good and 'very average' points in the objective space. The different behaviour of the two algorithms can be easily understood as NSGA-II has no means of stopping its exploitation (i.e. its repeated sampling) of a fit point, other than its niching

TABLE IX

MEDIAN AND INTERQUARTILE RANGE VALUES OF THE BINARY EPSILON INDICATOR AFTER 250/260 EVALUATIONS OF PAREGO/NSGA-II. LOWER VALUES INDICATE BETTER PERFORMANCE. THE DISTRIBUTIONS OF THE I VALUES ARE TESTED USING THE MANN-WHITNEY RANK SUM TEST. THE z VALUES AND SIGNIFICANCE LEVEL ARE INDICATED. PAREGO IS SIGNIFICANTLY BETTER THAN NSGA-II UNLESS STATED

Function	median values		IQR values		z-value	significance
	$(I_{ParEGO}, I_{NSGA-II})$	$(I_{ParEGO}, I_{NSGA-II})$	$(I_{ParEGO}, I_{NSGA-II})$	$(I_{ParEGO}, I_{NSGA-II})$		
KNO1	(1.2906, 1.7204)	(0.8487, 1.4558)	-2.3520	> 99%		
OKA1	(0.2674, 0.7440)	(0.1268, 0.1605)	-5.4211	> 99%		
OKA2	(-0.0001, 1.3058)	(0.0373, 0.3484)	-5.5468	> 99%		
VLMOP2	(0.0467, 0.0491)	(0.0205, 0.0784)	-1.043963	- no winner		
VLMOP3	(0.1801, 0.0866)	(0.1256, 0.2509)	-0.6666	- no winner		
DTLZ1a	(0.0817, 4.7435)	(0.1699, 8.5125)	-5.5468	> 99%		
DTLZ2a	(0.2425, 0.2141)	(0.1161, 0.3432)	-1.2452	- no winner		
DTLZ4a	(0.2674, 0.8682)	(0.3878, 0.2488)	-4.4400	> 99%		
DTLZ7a	(0.0152, 3.025)	(0.0516, 2.2221)	-5.5468	> 99%		

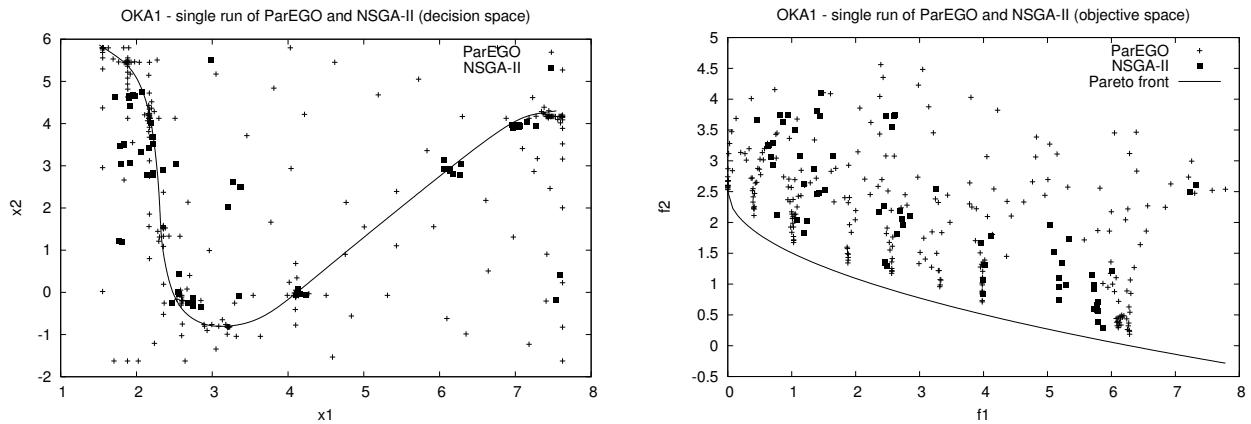


Fig. 10. Solutions in parameter space (left) with the Pareto set shown as a dark line. The corresponding points in objective space (right) for the first run of ParEGO and of NSGA-II after 250/260 function evaluations on the OKA1 test function. Notice how ParEGO explores the parameter space much more evenly and yet still finds more good points

policy, whereas ParEGO does not visit points that have little or no expectation of being improved (according to its landscape model), no matter how fit they currently are with respect to other solutions.

IX. SUMMARY AND CONCLUSION

In many optimization scenarios, the number of fitness evaluations that can be performed is severely limited by cost constraints. In this study, the performance of two multiobjective optimization algorithms, ParEGO and NSGA-II, was measured on much shorter runs than used in most previous MOEA studies. A suite of nine difficult, but low-dimensional, multiobjective test functions of limited ruggedness were used to evaluate the algorithms, and on these, it was confirmed that both ParEGO and NSGA-II outperformed a random search, even over a very small number of function evaluations.

ParEGO generally outperformed NSGA-II on the tested functions, at both 100 and 250 function evaluations, especially when the worst-case performance was measured. This suggests that ParEGO may offer a more effective search on problems like the instrument set-up optimization problem, described

herein, where only one function evaluation can be performed at a time. We are currently evaluating it on this problem as well as some other closely-related ‘wet experiment’ problems.

Although ParEGO’s performance was good on the tested functions, there are at least two potential drawbacks of the current version:

- 1) Normalization of the cost space relies on knowledge of the cost limits,
- 2) The use of uniformly random scalarizing vectors does not necessarily result in the best distribution of nondominated points. Some parts of the Pareto front may be far easier to find than others, so that this may lead to a poor approximation.

In future work, it is our intention to investigate a method based on adapting the scalarizing vectors, as a function of the points visited in the objective space, in order to obtain a better improvement to the current nondominated front at each iteration.

The addition of mechanisms to deal with integer and mixed-integer problems, as well as constraints is also under way.

Acknowledgments

Thanks to Douglas Kell for leading the collaborative work on the instrument set-up optimization problem that led to this paper, and to Julia Handl and the anonymous reviewers for helpful comments and suggestions. JK is supported by a David Phillips Fellowship from the Biotechnology and Biological Sciences Research Council (BBSRC), UK.

REFERENCES

- [1] G. E. P. Box and N. R. Draper. *Empirical Model Building and Response Surfaces*. John Wiley and Sons, New York, 1987.
- [2] M. Brown, W. Dunn, D. Ellis, R. Goodacre, J. Handl, J. Knowles, S. O'Hagan, I. Spasic, and D. Kell. A metabolome pipeline: from concept to data to knowledge. *Metabolomics*, 1(1):39–51, 2005.
- [3] D. Büche, G. Guidati, P. Stoll, and P. Kourmoursakos. Self-organizing maps for Pareto optimization of airfoils. In J. J. Merelo Guervós et al., editors, *Parallel Problem Solving from Nature—PPSN VII*, pages 122–131. Springer-Verlag, Lecture Notes in Computer Science Volume 2439, 2002.
- [4] D. Buche, N. Schraudolph, and P. Koumoutsakos. Accelerating evolutionary algorithms with Gaussian process fitness function models. *IEEE Transactions on Systems, Man, and Cybernetics C*, 34, 2004. To appear.
- [5] P. B. Cheung, L. F. Reis, K. T. Formiga, F. H. Chaudhry, and W. G. Ticona. Multiobjective evolutionary algorithms applied to the rehabilitation of a water distribution system: a comparative study. In C. M. Fonseca et al., editors, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 662–676. Springer, Lecture Notes in Computer Science Volume 2632, 2003.
- [6] C. A. Coello Coello. An updated survey of GA-based multiobjective optimization Techniques. *ACM Computing Surveys*, 32(2):109–143, 2000.
- [7] C. A. Coello Coello and G. Toscano Pulido. A micro-genetic algorithm for multiobjective optimization. In E. Zitzler et al., editors, *First International Conference on Evolutionary Multi-Criterion Optimization*, pages 126–140. Springer-Verlag, Lecture Notes in Computer Science Volume 1993, 2001.
- [8] C. A. Coello Coello, G. Toscano Pulido, and M. Salazar Lechuga. Handling multiple objectives with particle swarm optimization. *IEEE Transactions on Evolutionary Computation*, 8(3):256–279, 2004.
- [9] C. A. Coello Coello, D. A. Van Veldhuizen, and G. B. Lamont. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic Publishers, New York, 2002. ISBN 0-3064-6762-3.
- [10] D. A. Cohn, L. Atlas, and R. E. Ladner. Improving generalization with active learning. *Machine Learning*, 15(2):201–221, 1994.
- [11] D. Corne, M. Oates, and D. Kell. Landscape state machines: tools for evolutionary algorithm performance analyses and landscape/algorithm mapping. In *Applications of Evolutionary Computing*, pages 187–198. Springer, Lecture Notes in Computer Science Volume 2611, 2003.
- [12] D. W. Corne, N. R. Jerram, J. D. Knowles, and M. J. Oates. PESA-II: Region-based selection in evolutionary multiobjective optimization. In L. Spector et al., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2001)*, pages 283–290, San Francisco, California, 2001. Morgan Kaufmann Publishers.
- [13] V. Coverstone-Carroll, J. Hartmann, and W. Mason. Optimal multi-objective low-thrust spacecraft trajectories. *Computer Methods in Applied Mechanics and Engineering*, 186:387–402, 2000.
- [14] Z. S. Davies, R. J. Gilbert, R. J. Merry, D. B. Kell, M. K. Theodorou, and G. W. Griffith. Efficient improvement of silage additives by using genetic algorithms. *Applied and Environmental Microbiology*, pages 1435–1443, April 2000.
- [15] K. Deb. Evolutionary algorithms for multi-criterion optimization in engineering design. In K. Miettinen et al., editors, *Evolutionary Algorithms in Engineering and Computer Science*, chapter 8, pages 135–161. John Wiley & Sons, Ltd, Chichester, UK, 1999.
- [16] K. Deb. *Multi-Objective Optimization using Evolutionary Algorithms*. John Wiley & Sons, Chichester, UK, 2001. ISBN 0-471-87339-X.
- [17] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. KanGAL report 200001, Indian Institute of Technology, Kanpur, India, 2000.
- [18] K. Deb and H. Beyer. Self-adaptive genetic algorithms with simulated binary crossover. *Evolutionary Computation*, 9(2):197–221, 2001.
- [19] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable test problems for evolutionary multi-objective optimization. Technical Report 112, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, 2001.
- [20] M. Emmerich, N. Beume, and B. Naujoks. An EMO algorithm using the hypervolume measure as selection criterion. In *Evolutionary Multi-Criterion Optimization — EMO 2005*, volume 3410 of LNCS, pages 62–76. Springer, 2005.
- [21] M. Emmerich, A. Giotis, M. Özdemir, T. Bäck, and K. Giannakoglou. Metamodel-assisted evolution strategies. In *Parallel Problem Solving from Nature — PPSN VII*, volume 2439 of LNCS, pages 361–370, Berlin Heidelberg, 2002. Springer-Verlag.
- [22] M. Emmerich and B. Naujoks. Metamodel-assisted multiobjective optimisation strategies and their application in airfoil design. In I. Parmee, editor, *Adaptive Computing in Design and Manufacture VI*, pages 249–260, London, 2004. Springer.
- [23] M. Emmerich and B. Naujoks. Metamodel-assisted multiobjective optimisation strategies with implicit constraints and their application in airfoil design. In K. Giannakoglou et al., editor, *ERCOFTAC 2004*, page (CDROM). CIMNE, 2004.
- [24] J. R. G. Evans, M. J. Edirisinghe, and P. V. C. J. Eames. Combinatorial searches of inorganic materials using the inkjet printer: science philosophy and technology. *Journal of the European Ceramic Society*, 21:2291–2299, 2001.
- [25] M. Farina, K. Deb, and P. Amato. Dynamic multiobjective optimization problems: Test cases, approximations, and applications. *IEEE Transactions on Evolutionary Computation*, 8(5):425–442, 2004.
- [26] C. M. Fonseca and P. J. Fleming. On the performance assessment and comparison of stochastic multiobjective optimizers. In H.-M. Voigt, W. Ebeling, I. Rechenberg, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature—PPSN IV*, pages 584–593. Springer-Verlag, Lecture Notes in Computer Science, 1996.
- [27] A. Gaspar-Cunha and A. Vieira. A multi-objective evolutionary algorithm using approximate fitness evaluation. In G. Bugeđa, J. A. Dsidri, J. Periaux, M. Schoenauer, and G. Winter, editors, *International Congress on Evolutionary Methods for Design, Optimization and Control with Applications to Industrial Problems EUROGEN*, 2003.
- [28] A. Gaspar-Cunha and A. S. Vieira. A hybrid multi-objective evolutionary algorithm using an inverse neural network. Hybrid Metaheuristics (HM 2004) Workshop at ECAI 2004, Valencia, Spain, 2004.
- [29] K. C. Giannakoglou. Design of optimal aerodynamic shapes using stochastic optimization methods and computational intelligence. *International Review Journal Progress in Aerospace Sciences*, 38:43–76, 2002.
- [30] J. Horn and N. Nafpliotis. Multiobjective optimization using the niched Pareto genetic algorithm. Technical Report IlliGAI Report 93005, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA, 1993.
- [31] E. J. Hughes. Multi-objective binary search optimisation. In C. M. Fonseca, P. J. Fleming, E. Zitzler, K. Deb, and L. Thiele, editors, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 102–117. Springer, Lecture Notes in Computer Science Volume 2632, 2003.
- [32] M. T. Jensen. Reducing the run-time complexity of multiobjective EAs: The NSGA-II and other algorithms. *IEEE Transactions on Evolutionary Computation*, 7(5):503–515, 2003.
- [33] Y. Jin. A comprehensive survey of fitness approximation in evolutionary computation. *Soft Computing*, 9(1):3–12, 2005.
- [34] Y. Jin, M. Olhofer, and B. Sendhoff. A framework for evolutionary optimization with approximate fitness functions. *IEEE Transactions on Evolutionary Computation*, 6(5):481–494, 2002.
- [35] D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998.
- [36] J. Knowles. Parameter setting for PESA-II on the closed-loop spectrometer optimization problem: a simulation study using short run experiments. Internal report available on request from j.knowlesumist.ac.uk, April 2004.
- [37] J. Knowles. Supporting material for this paper. <http://dbk.ch.umist.ac.uk/knowles/parego/>, 2004.
- [38] J. Knowles and D. Corne. On metrics for comparing nondominated sets. In *Congress on Evolutionary Computation (CEC'2002)*, volume 1, pages 711–716, Piscataway, New Jersey, 2002. IEEE Service Center.
- [39] J. D. Knowles. *Local-Search and Hybrid Evolutionary Algorithms for Pareto Optimization*. PhD thesis, University of Reading, UK, 2002.
- [40] J. D. Knowles and D. W. Corne. Approximating the nondominated front using the Pareto archived evolution strategy. *Evolutionary Computation*, 8(2):149–172, 2000.

- [41] K.-H. Liang, X. Yao, and C. Newton. Evolutionary search of approximated n-dimensional landscapes. *International Journal of Knowledge-Based Intelligent Engineering Systems*, 4(3):172–183, 2000.
- [42] R. Myers and D. Montgomery. *Response Surface Methodology*. Wiley, New York, 1995.
- [43] P. K. S. Nain and K. Deb. A computationally effective multi-objective search and optimization technique using coarse-to-fine grain modeling. Technical Report Kangal Report No. 2002005, IITK, Kanpur, India, 2002.
- [44] B. Naujoks, L. Willmes, T. Bäck, and W. Haase. Evaluating multi-criteria evolutionary algorithms for airfoil optimization. In J. J. Merelo Guervós et al., editors, *Parallel Problem Solving from Nature—PPSN VII*, pages 841–850. Springer-Verlag, Lecture Notes in Computer Science Volume 2439, 2002.
- [45] S. O’Hagan, W. B. Dunn, M. Brown, J. D. Knowles, and D. B. Kell. Closed-loop, multiobjective optimization of analytical instrumentation: gas chromatography/time-of-flight mass spectrometry of the metabolomes of human serum and of yeast fermentations. *Analytical Chemistry*, 77(1):290–303, 2004.
- [46] T. Okabe, Y. Jin, M. Olhofer, and B. Sendhoff. On test functions for evolutionary multi-objective optimization. In *Proceedings of the Eighth International Conference on Parallel Problem Solving from Nature*, LNCS. Springer, 2004. (To appear).
- [47] Y. S. Ong, P. B. Nair, and A. J. Kean. Evolutionary optimization of computationally expensive problems via surrogate modeling. *AIAA Journal*, 41(4):687–696, 2003.
- [48] Y. S. Ong, P. B. Nair, A. J. Keane, and Z. Z. Zhou. Surrogate-assisted evolutionary optimization frameworks for high-fidelity engineering design problems. In Y. Jin, editor, *Knowledge Incorporation in Evolutionary Computation*, Studies in Fuzziness and Soft Computing Series. Springer Verlag, 2004.
- [49] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.
- [50] A. Ratle. Accelerating the convergence of evolutionary algorithms by fitness landscape approximation. In *Parallel Problem Solving from Nature - PPSN V*, pages 87–96, 1998.
- [51] R. G. Regis and C. A. Shoemaker. Local function approximation in evolutionary algorithms for the optimization of costly functions. *IEEE Transactions on Evolutionary Computation*, 8(5):490–505, 2004.
- [52] J. Sacks, W. Welch, T. Mitchell, and H. Wynn. Design and analysis of computer experiments (with discussion). *Statistical Science*, 4(409–435), 1989.
- [53] K. I. Smith, R. M. Everson, and J. E. Fieldsend. Dominance measures for multi-objective simulated annealing. In *2004 Congress on Evolutionary Computation (CEC’2004)*, volume 1, pages 23–30, Portland, Oregon, USA, 2004. IEEE Service Center.
- [54] S. Vaidyanathan, D. I. Broadhurst, D. B. Kell, and R. Goodacre. Explanatory optimization of protein mass spectrometry via genetic search. *Analytical Chemistry*, 75(23):6679–6686, 2003.
- [55] S. Vaidyanathan, D. Kell, and R. Goodacre. Selective detection of proteins in mixtures using electrospray ionization mass spectrometry: influence of instrumental settings and implications for proteomics. *Analytical Chemistry*, 76:5024–5032, 2004.
- [56] D. A. V. Veldhuizen and G. B. Lamont. Multiobjective evolutionary algorithm test suites. In J. Carroll et al., editors, *Proceedings of the 1999 ACM Symposium on Applied Computing*, pages 351–357, San Antonio, Texas, 1999. ACM.
- [57] A. Vicini and D. Quagliarella. Airfoil and wing design using hybrid optimization strategies. *AIAA Journal*, 37(5), 1999.
- [58] D. Weuster-Botz and C. Wandrey. Medium optimization by genetic algorithm for continuous production of formate dehydrogenase. *Process Biochemistry*, 30:563–571, 1995.
- [59] D. H. Wolpert. The lack of a priori distinctions between learning algorithms. *Neural Computation*, 8:1341–1390, 1996.
- [60] E. Zitzler. *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*. PhD thesis, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, 1999.
- [61] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, April 2003.



Joshua Knowles received the B.Sc. (Hons.) degree in Physics with Mathematics in 1993, the M.Sc. degree in Information Systems Engineering (at Distinction level) in 1996, and the Ph.D. in 2002, all from the University of Reading, UK.

From 2001–2003 he worked as a European commission Marie Curie research fellow at IRIDIA, Free University of Brussels, Belgium. Since 2003 he has held a BBSRC David Phillips fellowship at the University of Manchester, where his work is focused on biochemistry applications of heuristic optimization, machine learning and data-mining. He is the Program Chair of the 10th Online World Conference on Soft Computing in Industrial Applications (WSC10).