

Playing in the Objective Space: Coupled Approximators for Multi-Objective Optimization

Harold Soh*, Ong Yew Soon†, Mohamed Salahuddin*, and Terence Hung*

*Institute of High Performance Computing

1 Science Park Road, 01-01 The Capricorn

Singapore Science Park II

Singapore 117528

Email: sohsh@ihpc.a-star.edu.sg, mohdsh@ihpc.a-star.edu.sg,

terence@ihpc.a-star.edu.sg

†School of Computer Engineering

Nanyang Technological University

Blk N4, 2b-39, Nanyang Avenue

Singapore 639798

Email: asysong@ntu.edu.sg

Abstract— This paper presents a method of integrating computational intelligence and the operators used in evolutionary algorithms. We investigate approximation models of the objective function and its inverse and propose two simple algorithms that use these *coupled approximators* to optimize multi-objective functions. This method is a break from traditional approach used by standard cross-over and mutation operators, which only explore the objective space through “near-blind” manipulation of solutions in the parameter space. Fundamentally, our proposed *intelligent operators* use learned models of the coupling between the objective space and the parameter space to generate successively better solutions by extrapolating (or interpolating) from known solutions directly in the objective space. We term our implementation of the developed techniques as the Coupled Approximators Evolutionary Algorithm (CAEA). Empirical results of tests on scalable benchmark functions from the DTLZ test suite suggest that CAEA is an effective method for generating good solutions within a tight computational budget. These promising results prompt us to suggest several avenues for future research including combination with local search methods, incorporation of domain-knowledge and more efficient search algorithms.

I. INTRODUCTION

The Evolutionary Algorithms (EA) arena has shown tremendous activity in the past few years, with advancements being made on both the theoretical and applied fronts. One particular area of interest has been the development of multi-objective evolutionary methods that are capable of optimizing expensive and difficult problems quickly and reliably. Previous work in this area has focussed mainly on using inexpensive models of the exact evaluation function. This work explores an alternative approach by merging computational intelligence methods with evolutionary algorithms to derive *intelligent operators* which generate candidate solutions based on models of the objective function and the function’s inverse. We term the models developed as *coupled approximators* and present two algorithms that use these models to explore the objective space.

This paper is organized as follows: The section immedi-

ately following this introduction explores some relevant work on multi-objective optimization with evolutionary algorithms, fitness approximation and competent operators. Section III presents details of our proposed method, i.e., the Coupled Approximators Evolutionary Algorithm (CAEA). In Section IV, we present results and an analysis of simulations with scalable benchmark problems from the DTLZ [1] test suite. Section V concludes this paper with possible future work and conclusions derived from this study.

II. BACKGROUND

In this section, we cover related work in the fields of multi-objective optimization, evolutionary algorithms and fitness approximation.

A. Multi-objective Optimization with Evolutionary Algorithms

Intuitively, the problem of multi-objective optimization can be viewed as a search through a m -dimensional space for all minimum (or maximum) objective vectors that satisfy posed constraints. The search is usually performed “indirectly” by varying parameter vectors, also called decision vectors, in a d -dimensional space. From this point forward, we assume, without loss in generality, minimization problems.

We begin by defining m -dimensional fitness space of all feasible solutions. Given an objective function $f(\hat{x}) : \mathbb{R}^d \rightarrow \mathbb{R}^m$ with finite range and constraint functions $g(\hat{x}) : \mathbb{R}^d \rightarrow \mathbb{R}^k$ and $h(\hat{x}) : \mathbb{R}^d \rightarrow \mathbb{R}^l$, the feasible objective space is defined as the set of all vectors in the function’s range that satisfy the given constraints, $F = \{\hat{f}_i \in \mathbb{R}^m | \hat{f}_i = f(\hat{x}_i) \wedge (g(\hat{x}_i) > \mathbf{0}) \wedge (h(\hat{x}_i) = \mathbf{0})\}$.

Since our goal is to find the best solutions in a given F , it is helpful to consider what it means for one solution to be better than or *dominates* another. Consider $\hat{f} = (f_1, f_2, \dots, f_m)$, $\hat{g} = (g_1, g_2, \dots, g_m) \in F$. \hat{f} is said to dominate \hat{g} , denoted as $\hat{f} \succ \hat{g}$ iff $\forall i \in \{1, 2, \dots, m\} : f_i \leq g_i \wedge \exists j \in \{1, 2, \dots, m\} : f_j < g_j$.

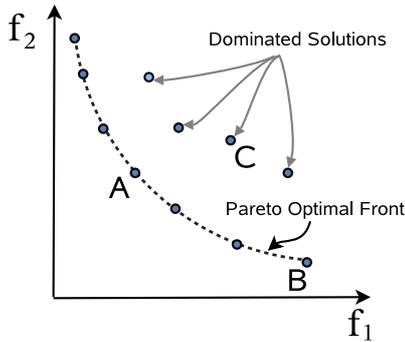


Fig. 1. A sample convex Pareto optimal front for a minimization problem.

For example, in fig 1, solution A dominates solution C while solutions B and C are non-dominating.

A Pareto optimal vector (or point), denoted as \hat{f}^* , is a vector which is not dominated by any other vector in the objective space ,i.e., $\nexists \hat{f}_j \in F : \hat{f}_j \succ \hat{f}^*$. The Pareto optimal set or front, defined as the set of all Pareto optimal vectors, is denoted as $F^* = \{\hat{f}_i^* \in F | \nexists \hat{f}_j \in F : \hat{f}_j \succ \hat{f}_i^*\}$. An example of a Pareto-optimal front for a minimization problem is illustrated in fig. 1.

Multi-objective optimization is defined as a search for the Pareto optimal front, F^* and the corresponding $X = \{\hat{x} \in \mathbb{R}^p | f(\hat{x}) \in F^*\}$. Multi-objective optimization problems are abundant in the real-world and include tasks such as scheduling, routing and structural design.

In recent years, effective evolutionary algorithms have been developed to solve multi-objective problems. The more influential of these multi-objective evolutionary algorithms (MOEA) include the Non-Dominating Sorting Genetic Algorithm-II (NSGA-II) [2], the Strength Pareto Evolutionary Algorithm 2 (SPEA2) [3] and the Pareto Archived Evolutionary Strategy (PAES) [4]. Readers wanting a comprehensive survey and history of such methods are directed to [5] and [6].

B. Evolutionary Algorithms with Fitness Approximation

A limitation of multi-objective evolutionary algorithms is that a relatively large number of evaluations are required to converge to the Pareto optimal front. This may prohibit the use of MOEAs when evaluations are expensive, for example simulating an engine to determine thermodynamic properties. Previous work attempting to resolve this problem involved using models to approximate the fitness function and used the models to evaluate solutions instead of the true objective function, for example [7] [8] [9] [10] [11] [12]. For more details on some of the evolutionary computational frameworks that employ approximation models, the reader is referred to the survey papers in [13] and [14].

C. Competent Genetic Algorithms

Another line of research has explored using models such as bayesian networks [15] and marginal product models [16]

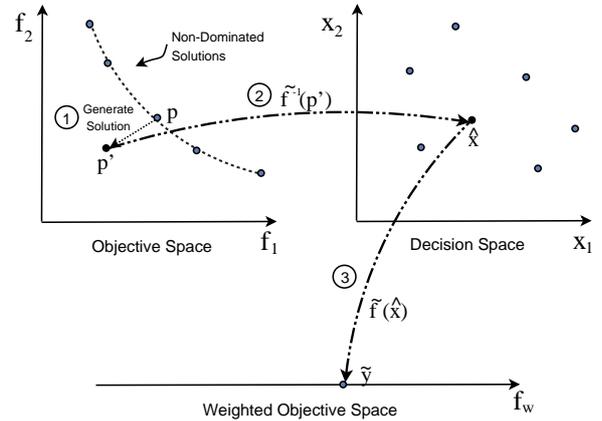


Fig. 2. Optimization using coupled approximators is a three-step process. First the solutions are generated by extrapolation or interpolation in the objective space from a non-dominated solution, p . Next, the inverse model is used to obtain the estimated decision vector, \hat{x} . This decision vector is then verified by the approximated objective function. If the resultant estimated weighted objective function is less than the weighted objective values of the p , the solution is accepted for evaluation by the true objective function.

as “competent” or “linkage-learning” operators. These operators use probabilistic models of genotype structures, *building blocks*, that give good fitness values and have been shown to solve certain types of problems more effectively. In this paper, we present an alternative approach of *intelligent operators* which, instead of learning linkages, utilizes learned models of the fitness space and its inverse.

III. COUPLED APPROXIMATORS FOR MULTI-OBJECTIVE OPTIMIZATION

The fundamental idea in this work is to model the objective function, its inverse and apply these models towards locating optimal solutions in the search space. We term these two models used in conjunction as *coupled approximators*. At each iteration of the optimization process, dominating or non-dominating solutions are constructed by extrapolating (or interpolating) from existing solutions in the objective space. The inverse model is used to obtain the estimated decision vector which is then verified by the approximated objective function. This three-step process is illustrated by fig. 2.

The Coupled Approximators Evolutionary Algorithm (CAEA) pseudocode (fig. 3) illustrates a conceptually simple algorithm based on the Non-dominated Sorting Algorithm (NSGA-II) [2]. The algorithm begins as a standard population based evolutionary algorithm by initializing a population of size N via random or latin hypercube sampling. Solutions are then evaluated and sorted according to non-domination ranks.

Next, the algorithm creates γN new individuals. The new operators are used first ,i.e., the forward and inverse approximation models are created and used to generate new solutions. It is possible that these operators fail to return valid solutions and if there is space remaining in the new population buffer, the remaining solutions are generated using standard crossover operators. The new solutions are evaluated and added to the current population. The merged population is

Coupled Approximators Evolutionary Algorithm, CAEA

- 1) $P_0 = \text{Initialize}(N)$
 - 2) $\text{Evaluate}(P_0)$
 - 3) $\text{NonDominatedSort}(P_0)$
 - 4) $t := 0$
 - 5) **while** termination conditions not met **do**
 - 6) $S_{CA} := \text{GenerateSolutionsCA}(P_t, t)$:
 - 7) $S_{SO} := \text{StandardOperators}(P_t, \gamma N - |S_{CA}|)$
 - 8) $S := S_{CA} \cup S_{SO}$
 - 9) $\text{Evaluate}(S)$
 - 10) $P' = P_t \cup S$
 - 11) $\text{NonDominatedSort}(P')$
 - 12) $P_{t+1} := \text{Truncate}(P_{t+1}, N)$
 - 13) $t := t + 1$
 - 14) **done**
-

Fig. 3. The Coupled Approximators Evolutionary Algorithm.

then sorted into non-domination ranks and truncated based on rank and crowding distance. The following subsections detail the modeling methods and solution construction algorithms, the `GenerateSolutionsCA` function (fig. 3, line 6).

A. Generating New Solutions

The `GenerateSolutionsCA` method first extracts the non-dominated front, P^* from the population and generates the coupled approximators by modeling fitness function with samples from the current population. As previously described, two models are generated:

- 1) The *forward model*, $\tilde{f}(\hat{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$, which is an approximation of a *weighted* objective function, $f_w(\hat{x}) = \sum_{i=1}^m w_i f_i(\hat{x})$. For simplicity, we consider equal weights in this study, i.e., $w_1 = w_2 = \dots = w_m = \frac{1}{m}$.
- 2) The *inverse model*, $\tilde{f}^{-1}(\hat{y}) : \mathbb{R}^m \rightarrow \mathbb{R}^d$, which is an approximation of the inverse objective function. In general, the inverse $f^{-1}(\hat{y})$ function may not exist so, the model approximates the inverse relation from the given samples.

Using P^* and the coupled approximators, dominating solutions (even iterations) or diverse, possibly non-dominating, solutions (odd iterations) are generated. The methods by which both types of solutions are constructed are as described in the following subsections.

B. Constructing Dominating Solutions

Recall that a given two solutions $\hat{f} = (f_1, f_2, \dots, f_m)$, $\hat{g} = (g_1, g_2, \dots, g_m) \in F$, \hat{f} dominates \hat{g} , $\hat{f} \succ \hat{g}$, iff $\forall i \in \{1, 2, \dots, m\} : f_i \leq g_i \wedge \exists j \in \{1, 2, \dots, m\} : f_j < g_j$. Hence, given a current population P , the algorithm (fig. 4) first extracts all non-dominated solutions, $P^* \subseteq P$ and for each non-dominated solution $\hat{p} \in P^*$, attempts to generate a dominating solution \hat{p}' by linearly extrapolating along each objective dimension, k .

GenerateSolutionsCA(P_t, t)

- 1) $P^* := \{p \in P_t \mid p \text{ is not dominated by any } p_i \in P_t\}$
 - 2) $\tilde{f}(\hat{x}) := \text{CreateForwardModel}(P_t)$.
 - 3) $\tilde{f}^{-1}(\hat{y}) := \text{CreateInverseModel}(P_t)$.
 - 4) **if** t is even **or** $t = 0$ **then**
 - 5) $S_{CA} := \text{GenDomingating}(P^*, \tilde{f}(\hat{x}), \tilde{f}^{-1}(\hat{y}), \alpha, \epsilon_k)$
 - 6) **else**
 - 7) $S_{CA} := \text{GenDiverse}(P^*, \tilde{f}(\hat{x}), \tilde{f}^{-1}(\hat{y}), \gamma N, \beta)$
 - 8) **return** S_{CA}
-

Fig. 4. Main method to generate dominating and non-dominated solutions relative to the current non-dominated front.

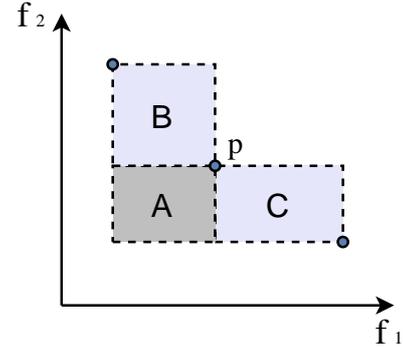


Fig. 5. Generated solution spaces for $\alpha = 1$ given some solution p . `GenDomingating` generates solutions within space A. `GenDiverse` generates solutions in spaces A, B and C.

The algorithm considers each objective dimension in turn and extends the point by some *step distance*, $d_k := \alpha \cdot (p_k - q_k)$ where \hat{q} is the next lowest value for the objective, i.e., \hat{q} s.t. $\nexists \hat{r} \in P^* : q_k < r_k < p_k$. The “aggression” parameter, α , determines how far along this objective to extend with $p_k - q_k$ as a unit. In fig. 5, the space A illustrates the neighbourhood region explored when $\alpha = 1$. A high α may cause a quick convergence to the front at the cost of making more errors. A smaller α is likely to yield a slower, more conservative search.

While not explicitly stated in the pseudocode, it is necessary to note that \hat{q} does not exist when p_k is the minimum. If $|P^*| > 1$, we can choose not to extrapolate along any objective that p_k is the minimum. Otherwise, d_k can be set to some predefined value or the difference between p_k and some minimum reference objective value, r_k . If the objective space is in \mathbb{R}^+ , r_k can be defaulted to 0.

The inverse model is then used to estimate the decision variables, \hat{x} , that would yield the objective vector \hat{p}' . The decision variables are then validated with the forward model. If the estimated weighted objectives, \tilde{y} is smaller than the weighted objective value of \hat{p} , the algorithm moves on to the next objective. Otherwise, the step distance is halved and the extension process repeated. The algorithm terminates if d_k falls below some predefined value, ϵ , for this objective. In this study, we found it convenient to set ϵ to be a percentage of the initial d_k value.

GenDominating(P^* , $\tilde{f}(\hat{x})$, $\tilde{f}^{-1}(\hat{y})$, α , ϵ)

```
1) for each solution  $\hat{p} \in P^*$  do
2)    $\hat{p}' := \hat{p}$ 
3)   for each objective  $k := 1, 2, \dots, m$  do
4)     Let  $\hat{q}$  s.t.  $\nexists \hat{r} \in P^* : q_k < r_k < p_k$ 
5)      $d_k := \alpha \cdot (p_k - q_k)$ 
6)     do
7)        $p'_k := p'_k - d_k$ 
8)        $\hat{x} := \tilde{f}^{-1}(\hat{p}')$ 
9)        $\hat{y} := \tilde{f}(\hat{x})$ 
10)      if  $\hat{y} < f_w(p)$ 
11)        end while loop
12)      else
13)         $d_k := d_k \div 2$ 
14)         $p'_k := p_k$ 
15)      while  $d_k > \epsilon$ 
16)    done
17)     $S_{CA} := S_{CA} \cup \hat{x}$ 
18) done
19) return  $S_{CA}$ 
```

Fig. 6. Algorithm to construct dominating solutions given a set of solutions, coupled approximation models, an “aggression” parameter, α and a termination condition, ϵ .

C. Constructing Diverse Solutions

This solution construction method (fig. 7) attempts to generate possibly non-dominating solutions to maintain a diverse pareto optimal front. The algorithm attempts to generate $\frac{\gamma N}{|P^*|}$ solutions from each non-dominated solution, $\hat{p} \in P^*$. A new solution p' is constructed by randomly perturbing the solution along each objective, k , within the range of the next smaller objective value, l_k , and the next larger objective value, u_k . Note if p_k is the minimum value, p_k is perturbed within some predefined range. As when constructing dominating solution, the inverse model is then used to estimate the corresponding decision variables, \hat{x} . Solutions are accepted for evaluation only if they are in the dominating or non-dominating space and are validated by the forward model. The spaces A, B and C in fig. 5 illustrate the neighbourhood region in the objective space that the GenDiverse algorithm explores.

IV. EMPIRICAL RESULTS

In this section, we present results on simulations of our modeling operators on several real-valued continuous test problems and compare the solutions found to those obtained via standard cross-over and mutation operators only.

A. Implementation and Experimental Setup

CAEA was implemented in C++ and utilized the Fast Artificial Neural Network Library (FANN) [17]. To model the objective function, $f(\hat{x})$ and its inverse, $f^{-1}(\hat{y})$, we used simple radial basis function (RBF) networks with $5 \times p$ internal nodes for the forward model and $5 \times (m + p)$ internal nodes for the inverse model.

GenDiverse(P^* , $\tilde{f}(\hat{x})$, $\tilde{f}^{-1}(\hat{y})$, γN , β)

```
1) for each solution  $\hat{p} \in P^*$  do
2)   for  $i := 1$  to  $\gamma N / |P^*|$  do
3)      $b := 0$ 
4)     do
5)        $\hat{p}' := \hat{p}$ 
6)       for each objective  $k := 1, 2, \dots, m$  do
7)         Let  $\hat{l}$  s.t.  $\nexists \hat{r} \in P^* : l_k < r_k < p_k$ 
8)         Let  $\hat{u}$  s.t.  $\nexists \hat{v} \in P^* : u_k > v_k > p_k$ 
9)          $p'_k = \text{random}(l_k, u_k)$ 
10)      done
11)       $\hat{x} := \tilde{f}^{-1}(\hat{p}')$ 
12)       $\hat{y} := \tilde{f}(\hat{x})$ 
13)       $b := b + 1$ 
14)      while  $(p' \neq p) \vee (\hat{y} > f_w(p)) \vee (b < \beta)$ 
15)       $S_{CA} = S_{CA} \cup \hat{p}'$ 
16)    done
17) done
18) return  $S_{CA}$ 
```

Fig. 7. Algorithm to construct non-dominated solutions given a set of solutions, coupled approximation models, the number of solutions to generate, γN and a termination condition, β .

TABLE I
PARAMETERS FOR CAEA AND NSGA-II

Parameter	CAEA	NSGA-II
Crossover Probability, p_c	-	0.9
Crossover Spread Factor, n_c	15	15
Mutation Probability, p_m	0.01	0.01
Mutation Spread Factor, n_m	20	20
Aggression, α	10	-
GenDominating Stopping criteria, ϵ	$0.1 \times \text{initial}d_k$	-
GenDiverse Stopping criteria, β	20	-
Solution multiplier, γ	1	-

We used the PISA [18] system as an experimental environment running on a 16 CPU Intel XEON 2.6Ghz shared memory system with 64 gigabytes of memory. CAEA was tested on the DTLZ2, DTLZ3, DTLZ6, DTLZ7 test problems [1] with 2, 4 and 8 objectives. The population sizes were set to 200, 300 and 400 for 2, 4 and 8 objectives respectively. The number of variables was set to the recommended $m + 9$ where m is the number of objectives. Each test was repeated 25 times with matched populations between algorithms. We compared CAEA against an algorithm using standard cross-over and mutation operators, NSGA-II, using the additive epsilon indicator, ϵ_+ and computed the Kruskal-Wallis statistical test [19]. The parameters for each algorithm is given in table I.

B. Results and Analysis

Figures 8 to 15 illustrate box plots of the additive epsilon indicator on the DTLZ test problems and tables II and III give the results of the Kruskal-Wallis statistical tests. We analyze snapshots of the runs at generations 25 and 50. At generation 25, the coupled approximators and construction algorithms generate significantly better solutions than those by standard

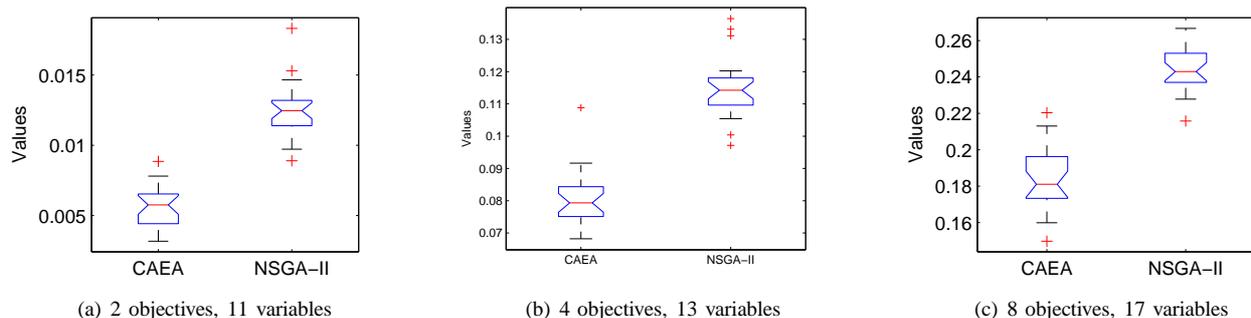


Fig. 8. ϵ_+ indicator values for CAEA and NSGA-II for DTLZ2 at Generation 25. Visually, CAEA outperforms NSGA-II for all three objective levels. This conclusion is confirmed by the Kruskal-Wallis statistical test.

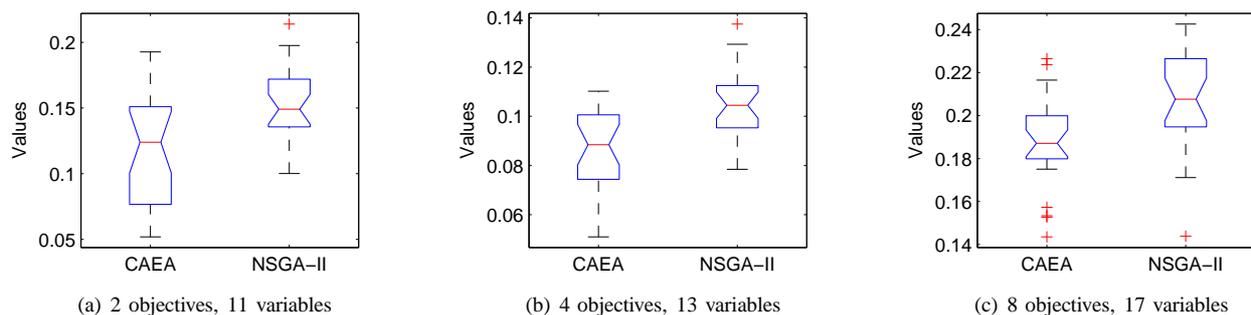


Fig. 9. ϵ_+ indicator values for CAEA and NSGA-II for DTLZ3 at Generation 25. As with DTLZ2, CAEA outperforms NSGA-II at all three objective values.

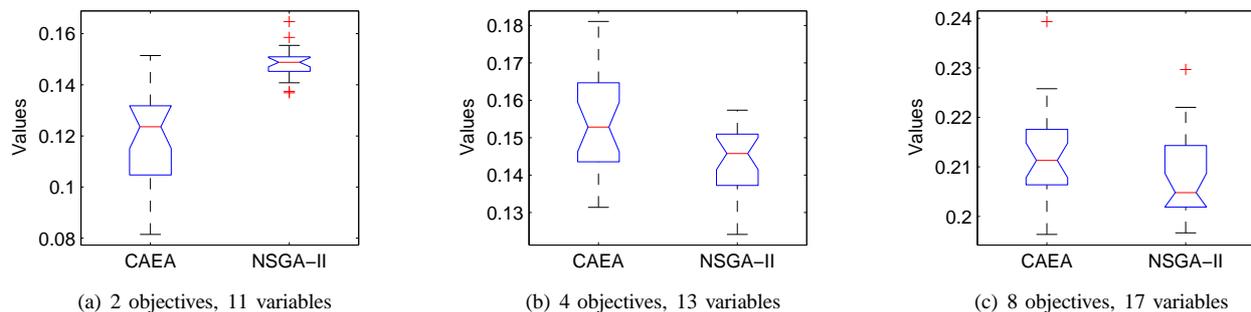


Fig. 10. ϵ_+ indicator values for CAEA and NSGA-II for DTLZ3 at Generation 25. CAEA is only able to outperform NSGA-II's standard operators on DTLZ6 with 2 objectives. No statistically significant results can be derived from the runs on higher objectives. However, a visual comparison indicates that CAEA performs poorer than NSGA-II on the higher objectives.

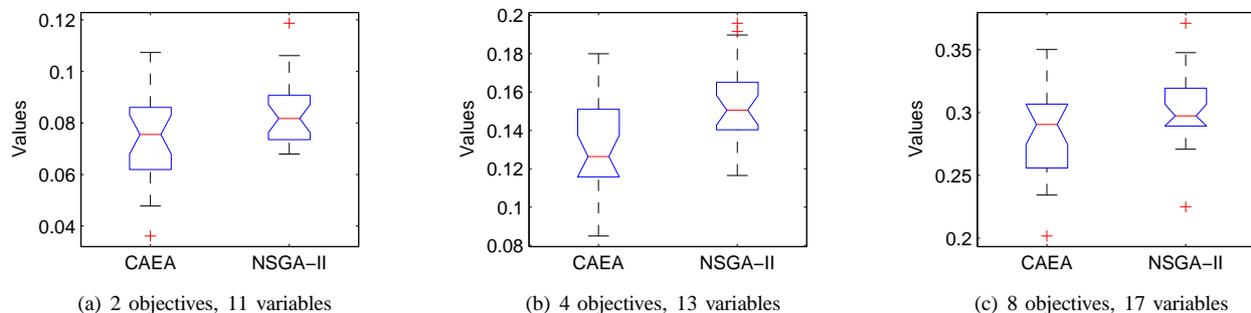


Fig. 11. ϵ_+ indicator values for CAEA and NSGA-II for DTLZ7 at Generation 25. CAEA is only able to outperform NSGA-II on DTLZ7 with 4 objectives. No statistically significant results can be derived from the runs on 2 and 8 objectives. However, a visual comparison suggests that CAEA performs better than NSGA-II even on higher objectives.

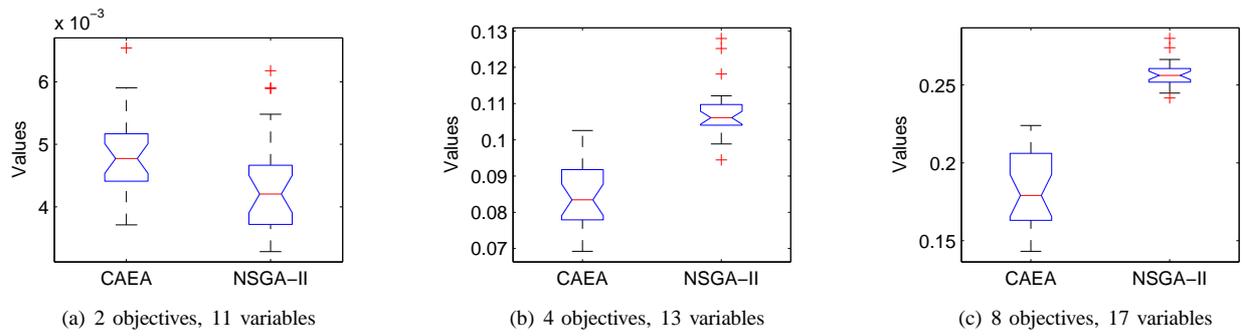


Fig. 12. ϵ_+ indicator values for CAEA and NSGA-II for DTLZ2 at Generation 50. CAEA performs better than NSGA-II on the problems with larger objectives. On DTLZ2 with 2 objectives, no statistically significant result can be concluded. However, a visual comparison suggests that NSGA-II has outperformed CAEA by the 50th generation.

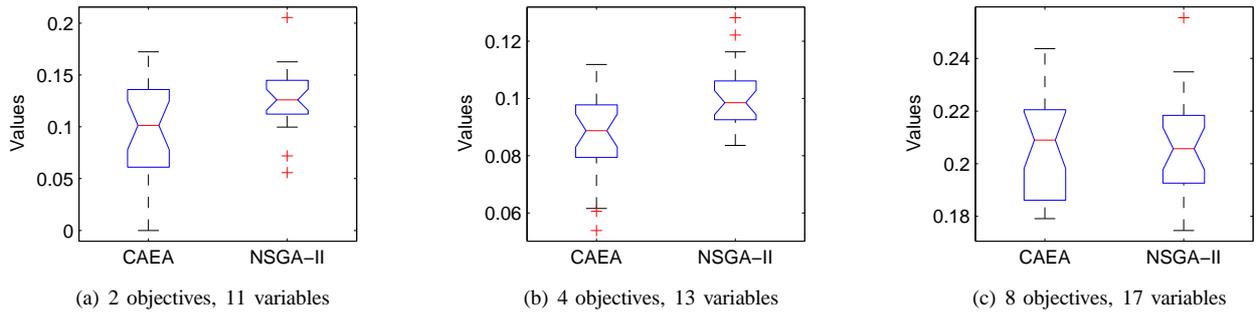


Fig. 13. ϵ_+ indicator values for CAEA and NSGA-II for DTLZ3 at Generation 50. No statistically significant result can be concluded for DTLZ3 with 2 and 8 objectives. With 4 objectives, CAEA outperforms NSGA-II.

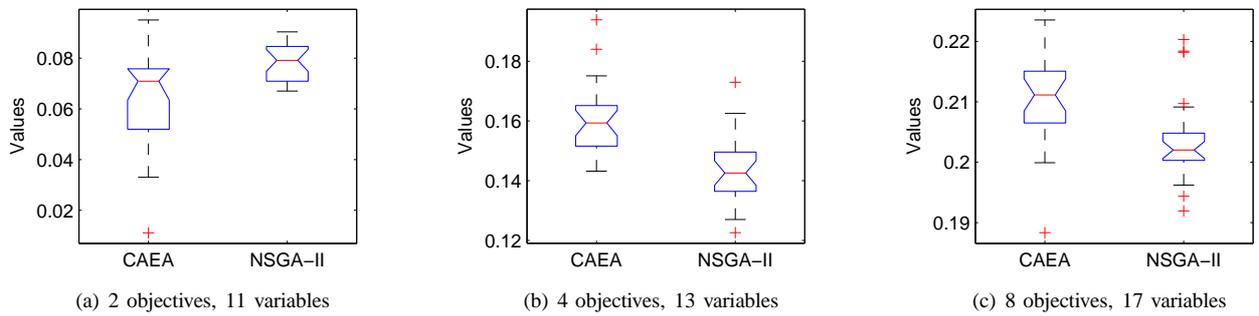


Fig. 14. ϵ_+ indicator values for CAEA and NSGA-II for DTLZ6 at Generation 50. CAEA outperforms NSGA-II on DTLZ6 with 2 objectives but performs poorly in comparison on the higher level objectives.

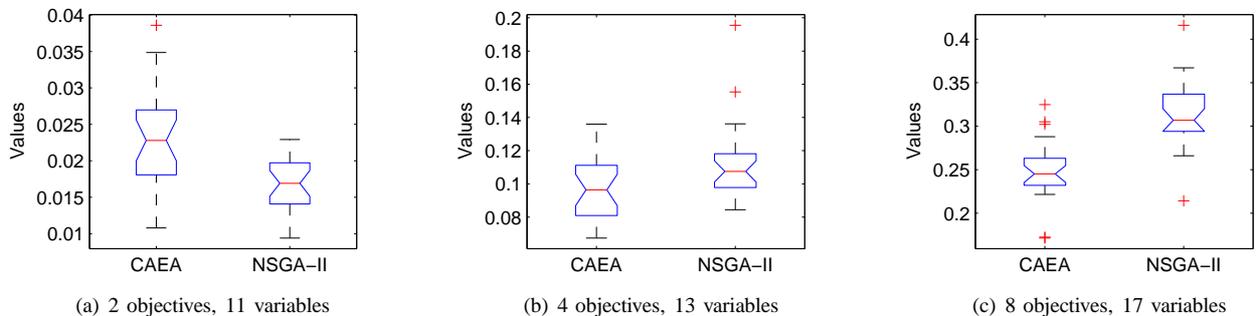
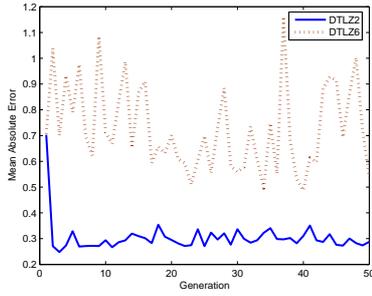
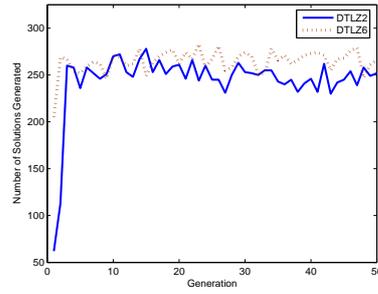


Fig. 15. ϵ_+ indicator values for CAEA and NSGA-II for DTLZ7 at Generation 50. CAEA finds significantly better solutions on DTLZ7 with 8 objectives. No statistically significant results can be obtained from the runs on 2 and 4 objectives. However, visual comparisons indicate that CAEA performs poorer compared to NSGA-II on the 2 objective problem and better on the 4 objective problem.



(a) Errors on DTLZ2 and DTLZ6 (4 objs)



(b) Number of Solutions Generated by Models

Fig. 16. While approximately the same number of solutions are generated on both problems, the coupled approximators makes significantly more errors on the DTLZ6 problem (dashed line) as compared to DTLZ2 (solid line).

TABLE II

KRUSKAL-WALLIS TEST FOR DTLZ PROBLEMS, GENERATION 25

	DTLZ2		
	$m = 2$	$m = 4$	$m = 8$
CAEA is better	$2.3e^{-16}$	$2.1e^{-15}$	$2.2e^{-16}$
NSGA-II is better	1.0	1.0	1.0
	DTLZ3		
	$m = 2$	$m = 4$	$m = 8$
CAEA is better	0.0030	$8.9e^{-05}$	0.0018
NSGA-II is better	0.99	0.99	0.99
	DTLZ6		
	$m = 2$	$m = 4$	$m = 8$
CAEA is better	$1.2e^{-10}$	H0	H0
NSGA-II is better	1.0	H0	H0
	DTLZ7		
	$m = 2$	$m = 4$	$m = 8$
CAEA is better	H0	0.00071	H0
NSGA-II is better	H0	0.99	H0

TABLE III

KRUSKAL-WALLIS TEST FOR DTLZ PROBLEMS, GENERATION 50

	DTLZ2		
	$m = 2$	$m = 4$	$m = 8$
CAEA is better	H0	$1.34e^{-15}$	$2.3e^{-16}$
NSGA-II is better	H0	1.0	1.0
	DTLZ3		
	$m = 2$	$m = 4$	$m = 8$
CAEA is better	H0	0.00040	H0
NSGA-II is better	H0	0.99	H0
	DTLZ6		
	$m = 2$	$m = 4$	$m = 8$
CAEA is better	0.0015	0.99	0.99
NSGA-II is better	0.99	$1.06e^{-06}$	0.00051
	DTLZ7		
	$m = 2$	$m = 4$	$m = 8$
CAEA is better	H0	H0	$1.61e^{-08}$
NSGA-II is better	H0	H0	1.0

evolutionary operators on DTLZ2, DTLZ3 and DTLZ7 with 2, 4 and 8 objectives.

The exception is DTLZ6, which is a deceptive problem and may have proved difficult for the RBFs to model accurately in higher dimensions. The models would then have made erroneous predictions and hence, generated poorer solutions. Evidence for this is suggested by fig. 16 and tbl. IV which illustrate the mean absolute errors made during a representative sample run on DTLZ6 as compared to DTLZ2.

At generation 50, CAEA outperforms NSGA-II on DTLZ2 (4 and 8 objectives), DTLZ3 (4 objectives), DTLZ6 (2 objectives) and DTLZ7 (8 objectives). For the rest of the problem instances, no statistical conclusion can be drawn except on the DTLZ6 problem (4 and 8 objectives) where NSGA-II outperforms CAEA.

The results indicate that for 3 of the 12 problems, i.e., DTLZ2 (2 Objectives), DTLZ3 (8 objectives) and DTLZ7 (2

TABLE IV

MEAN ERRORS ON DTLZ2 AND DTLZ6

Problem	Mean Error	Standard Deviation
DTLZ2, 4 Objs.	0.3036	0.0624
DTLZ6, 4 Objs.	0.7283	0.1683
Problem	Mean No. of Solutions Per Iteration	Standard Deviation
DTLZ2, 4 Objs.	244.78	33.5356
DTLZ6, 4 Objs.	264.48	12.4904

objectives), CAEA had lost its early lead, allowing standard operators to “catch-up”. Evidently, CAEA generated better solutions early in the optimization process and performance degraded later in the run. A plausible explanation for this phenomena is that as the population moved towards the Pareto optimal front, the samples became biased and the models were less able to generalize. To circumvent this problem in future studies, we propose the use of an archive to store previous

TABLE V
AVERAGE COMPUTATION TIME ON DTLZ2 IN SECONDS

Algorithm	2 Objectives	4 Objectives	8 Objectives
CAEA	722.4	2018.2	7133.5
NSGA-II	1.4	2.8	7.2

solutions from which to base our models and/or local learning.

C. Computational Time

The benefit obtained from using the coupled approximators is balanced by the extra computational time needed to create the models. Table V contrasts the computation time required by both algorithms on the DTLZ2 problem with 2, 4, and 8 averaged over 5 runs. The difference in computational time is significant (approximately 3 orders larger), hence making coupled-approximators best suited to cases where the function evaluations are relatively expensive. It may also be possible to reduce the computational time of CAEA by using cheaper approximation techniques such as lazy methods [20].

V. FUTURE WORK AND CONCLUSIONS

The experimental results obtained suggest that the use of models as operators, in particular the coupled approximators method, are promising alternatives or supplements to standard operators. On three of the four test problems (DTLZ2, DTLZ3 and DTLZ7), CAEA generated a better set of solutions early in the optimization process. However, these operators are not without bias as illustrated by the relatively poor results when tested against the DTLZ6 problem (4 and 8 objectives). It may be possible to improve on the results obtained on this deceptive problem by modifying the aggression parameter, α or by using alternative models or search methods. Following this line of thought opens up several avenues for research.

A. Dynamic Aggressiveness

The GenDominate function actively searches the objective space in an area determined by the aggression parameter, α . It is possible for α to be dynamically modified based on the errors made by the models during the optimization run ,i.e., α can be decreased when the models are observed to be making significant errors and increased when accuracy improves.

B. Clustering, Local Models and Archives

Prior research [21] [22] [12] suggests that the coupled approximators method may be improved by clustering the objective space and generating local models instead of global ones. The use of an archive to store previously evaluated individuals may also assist in the modeling process.

C. Incorporation of Domain Knowledge

The No Free Lunch (NFL) theorem [23] justifies the use of prior knowledge in the search and optimization of a particular problem. CAEA can be modified in the following two ways to utilize available prior information:

1) *Alternative Model Approximation Methods*: The implementation of CAEA developed used Radial Basis Function Networks to approximate the objective function and its inverse. However, a problem specific approximation or modeling technique can be used in its place. The choice of any one particular technique can be based on prior knowledge, required accuracy and computational time.

2) *Domain-specific Solution Construction Algorithms*: The solution construction algorithms presented in this paper are simple and were developed to illustrate the viability of using coupled approximators. It remains future research to develop more efficient general construction algorithms. If enough is known about the objective space, domain-specific algorithms capable of searching the space more efficiently can be used instead.

D. Combination with Local Search

Our empirical results suggest that the use of models as operators are more efficient than standard EA operators at generating good solutions early in the optimization process. A logical next step would be to combine these operators with local search algorithms. Our preliminary tests utilizing the Nelder Mead algorithm [24], sometimes referred to as the simplex algorithm, indicate that such hybrid algorithms yield significantly faster convergence.

ACKNOWLEDGMENT

This work was funded by the A*STAR SERC Grant (No. 052 015 0024) for the Grid-based PSE for Engineering of Materials (GPEM) project. The authors would like to thank Dr. Gary Lee Khee Khoo (IHPC) for his assistance in generating the box-plots presented in this paper.

REFERENCES

- [1] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, "Scalable multiobjective optimization test problems," in *Congress on Evolutionary Computation (CEC'2002)*, vol. 1, 2002, pp. 825–830. [Online]. Available: citeseer.ist.psu.edu/deb02scalable.html
- [2] K. Deb, S. Agrawal, A. Pratab, and T. Meyarivan, "A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II," in *Proceedings of the Parallel Problem Solving from Nature VI Conference*, M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. J. Merelo, and H.-P. Schwefel, Eds. Paris, France: Springer, Lecture Notes in Computer Science No. 1917, 2000, pp. 849–858. [Online]. Available: citeseer.ist.psu.edu/deb00fast.html
- [3] E. Zitzler, M. Laumanns, and L. Thiele, "SPEA2: Improving the Strength Pareto Evolutionary Algorithm," Gloriestrasse 35, CH-8092 Zurich, Switzerland, Tech. Rep. 103, 2001. [Online]. Available: citeseer.ist.psu.edu/zitzler02spea.html
- [4] J. D. Knowles and D. Corne, "Approximating the nondominated front using the pareto archived evolution strategy," *Evolutionary Computation*, vol. 8, no. 2, pp. 149–172, 2000. [Online]. Available: citeseer.ist.psu.edu/article/knowles00approximating.html
- [5] C. A. C. Coello, D. A. V. Veldhuizen, and G. B. Lamont, *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic Publishers, New York, 2002.
- [6] K. Deb, *Multi-Objective Optimization using Evolutionary Algorithms*. John Wiley and Sons, Chichester, UK, 2001.
- [7] J. Knowles, "ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems," *IEEE Transactions on Evolutionary Computation*, vol. 10, pp. 50–66, 2006.

- [8] Y. S. Ong, P. B. Nair, and A. J. Keane, "Evolutionary optimization of computationally expensive problems via surrogate modeling," *American Institute of Aeronautics and Astronautics Journal*, vol. 41, pp. 689–696, 2003. [Online]. Available: <http://ntu-cg.ntu.edu.sg/ysong/journal/GLSurrogate.pdf>
- [9] M. Emmerich, K. Giannakoglou, and B. Naujoks, "Single- and multi-objective evolutionary optimization assisted by gaussian random field metamodels." [Online]. Available: citeseer.ist.psu.edu/739286.html
- [10] A. Ratle, *Parallel Problem Solving from Nature - PPSN V*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 1998, vol. 1498/1998, ch. Accelerating the Convergence of Evolutionary Algorithms by Fitness Landscape Approximation.
- [11] Y. S. Ong, P. B. Nair, and K. Y. Lum, "Max-Min surrogate-Assisted evolutionary algorithm for robust aerodynamic design," *IEEE Transactions on Evolutionary Computation*, vol. 10, pp. 392–404, 2006.
- [12] Z. Z. Zhou, Y. S. Ong, P. B. Nair, A. J. Keane, and K. Y. Lum, "Combining global and local surrogate models to accelerate evolutionary optimization," *IEEE Transactions On Systems, Man and Cybernetics - Part C*, vol. 37, 2007. [Online]. Available: <http://ntu-cg.ntu.edu.sg/ysong/journal/GLSurrogate.pdf>
- [13] Y. Jin, "A comprehensive survey of fitness approximation in evolutionary computation," *Soft Computing Journal*, 2003. [Online]. Available: citeseer.ist.psu.edu/jin03comprehensive.html
- [14] Y. S. Ong, P. B. Nair, A. J. Keane, and K. W. Wong, *Surrogate-assisted evolutionary optimization frameworks for high-fidelity engineering design problems*. Springer Verlag, 2004, vol. 10, pp. 307–332.
- [15] M. Pelikan, K. Sastry, and D. E. Goldberg, "Multiobjective hboa, clustering, and scalability," in *GECCO '05: Proceedings of the 2005 conference on Genetic and evolutionary computation*. New York, NY, USA: ACM Press, 2005, pp. 663–670.
- [16] G. R. Harik, F. G. Lobo, and D. E. Goldberg, "The compact genetic algorithm," *IEEE-EC*, vol. 3, no. 4, p. 287, November 1999. [Online]. Available: citeseer.ist.psu.edu/harik98compact.html
- [17] S. Nissen. Fast artificial neural network library v2.0.0. [Online]. Available: <http://leenissen.dk/fann/>
- [18] S. Bleuler, M. Laumanns, L. Thiele, and E. Zitzler, "PISA — a platform and programming language independent interface for search algorithms," in *Evolutionary Multi-Criterion Optimization (EMO 2003)*, ser. Lecture Notes in Computer Science, C. M. Fonseca, P. J. Fleming, E. Zitzler, K. Deb, and L. Thiele, Eds. Berlin: Springer, 2003, pp. 494 – 508.
- [19] J. Knowles, L. Thiele, and E. Zitzler, "A tutorial on the performance assessment of stochastic multiobjective optimizers," Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich," 214, July 2005.
- [20] D. Aha, "Special issue on lazy learning," *Artificial Intelligence Review*, vol. 11, pp. 1–6, 1997.
- [21] H. Soh and M. Kirley, "moPGA: Towards a new generation of multi-objective genetic algorithms," in *Congress on Evolutionary Computation (CEC'2006)*, 2006, pp. 1702–1709.
- [22] R. Regis and C. Shoemaker, "Local function approximation in evolutionary algorithms for the optimization of costly functions," *IEEE Transactions on Evolutionary Computation*, vol. 8, pp. 490–505, 2004.
- [23] D. Wolpert and W. Macready, "No free lunch theorems for optimization," *IEEE Transactions on Evolutionary Computation*, vol. 1, pp. 67–82, 1997.
- [24] J. Nelder and R. Mead, "A simplex method for function minimization," *Computer Journal*, vol. 7, pp. 308–313, 1965.