Adaptive Reference Vector Generation for Inverse Model Based Evolutionary Multiobjective Optimization with Degenerate and Disconnected Pareto Fronts

Ran Cheng[1], Yaochu Jin[1,3], and Kaname Narukawa[2]

[1] Department of Computing, University of Surrey
Guildford, Surrey, GU2 7XH, United Kingdom
[2] Honda Research Institute Europe GmbH, 63073 Offenbach am Main, Germany
[3] College of Information Sciences and Technology, Donghua University
Shanghai 201620, China
{r.cheng,yaochu.jin}@surrey.ac.uk

Abstract. Inverse model based multiobjective evolutionary algorithm aims to sample candidate solutions directly in the objective space, which makes it easier to control the diversity of non-dominated solutions in multiobjective optimization. To facilitate the process of inverse modeling, the objective space is partitioned into several subregions by predefining a set of reference vectors. In the previous work, the reference vectors are uniformly distributed in the objective space. Uniformly distributed reference vectors, however, may not be efficient for problems that have nonuniform or disconnected Pareto fronts. To address this issue, an adaptive reference vector generation strategy is proposed in this work. The basic idea of the proposed strategy is to adaptively adjust the reference vectors according to the distribution of the candidate solutions in the objective space. The proposed strategy consists of two phases in the search procedure. In the first phase, the adaptive strategy promotes the population diversity for better exploration, while in the second phase, the strategy focused on convergence for better exploitation. To assess the performance of the proposed strategy, empirical simulations are carried out on two DTLZ benchmark problems, namely, DTLZ5 and DTLZ7, which have a degenerate and a disconnected Pareto front, respectively. Our results show that the proposed adaptive reference vector strategy is promising in tacking multi-objective optimization problems whose Pareto front is disconnected.

Keywords: Multiobjective Optimization, Model Based Evolutionary Optimization, Inverse Modeling, Reference Vectors

1 Introduction

A multiobjective optimization problem (MOP) involves several conflicting objectives to be optimized simultaneously. Without loss of generality, an MOP can be formulated as follows:

$$\begin{aligned} \min \quad & \boldsymbol{f}(\boldsymbol{x}) = (f_1(\boldsymbol{x}), f_2(\boldsymbol{x}), ..., f_m(\boldsymbol{x})) \\ \text{s.t.} \quad & \boldsymbol{x} \in X, \quad \boldsymbol{f} \in Y \end{aligned} \tag{1}$$

where $X \subset \mathbb{R}^n$ is the decision space and $\boldsymbol{x} = (x_1, x_2, ..., x_n) \in X$ is the decision vector, $Y \subset \mathbb{R}^m$ is the objective space and $\boldsymbol{f} \in Y$ is the objective vector, which is composed of m objective functions $f_1(\boldsymbol{x})$, $f_2(\boldsymbol{x})$,...,$f_m(\boldsymbol{x})$ that map $\boldsymbol{x}$ from X to Y. Due to the conflicting nature of the objectives, it is impossible to optimize all the objectives with one single solution. Consequently, there exists a set of optimal solutions, termed as Pareto optimal solutions, that trade-off between different objectives. The Pareto optimal solutions are often called the Pareto set in the decision space and image formed by the Pareto optimal solutions in the objective is termed Pareto front.

To obtain the Pareto optimal solutions, various multiobjective evolutionary algorithms (MOEAs) have been proposed, e.g. the elitist non-dominated sorting algorithm, known as NSGA-II [5], the decomposition based algorithm, called MOEA/D [14], among many others [16]. Most traditional MOEAs often require a high degree of diversity in storing the non-dominated solutions found so far in the current population or in an external archive. By contrast, model-based MOEAs [11, 12, 15] can alleviate the requirement on solution diversity by focusing on the construction of a probabilistic model in the decision space during the search. Such model based MOEAs, however, still rely on the use of a solution set, such as an archive, to represent the obtained non-dominated solutions. Another line of research that aims to alleviate the requirement on diversity is to build a regression model to represent the solutions obtained in the final generation by the optimizer [7, 9], which can be used to generate new solutions after the optimization process is complete, thereby enhancing the diversity of the final solutions. Inspired by the ideas in this line of research, a multiobjective evolutionary algorithm using Gaussian process based inverse modeling (IM-MOEA) has been proposed [2].

In IM-MOEA, an inverse model that maps candidate solutions in the objective space onto the decision space is built during the optimization. To facilitate the inverse modeling, the objective space is partitioned into several subregions using predefined *reference vectors*. By associating each candidate solution with a particular reference vector, a number of inverse models are built for each subregion by using the candidate solutions relating to this subregion as training data. In the previous work of IM-MOEA, the reference vectors are uniformly generated by means of the canonical simplex-lattice design method [3]. This method for generating reference vectors works well for MOPs with a continuous and uniform Pareto front. However, for some MOPs with a nonuniform or disconnected Pareto front, the predefined, uniformly distributed reference vectors may result in low efficiency, as some reference vectors may not be associated with any candidate solutions, thus causing a waste of computational resource. To tackle this problem, here we present an adaptive reference vector generation strategy for IM-MOEA. The proposed strategy is able to adapt the distribution of the reference vectors to the distribution of the candidate solutions in the objective space.

In the following, we first briefly introduce the recently proposed IM-MOEA in Section 2. Then the adaptive reference vector generation strategy is described in Section 3. Section 4 presents experimental results for assessing the performance of the proposed adaptive strategy. Finally, conclusion is drawn in Section 5.

2 IM-MOEA

Traditional EDAs aim to estimate the distribution of the candidate solutions in the decision space, while the models in IM-MOEA are built to represent the inverse mapping from the objective space to the decision space. With the inverse models thus built, evenly distributed candidate solutions can be directly sampled in the objective space and then mapped onto the decision space.

Considering that the estimation of the entire inverse mapping from the m-dimensional objective space to the n-dimensional decision space can be technically difficult, the multivariate inverse model is decomposed into a number of univariate regression models:

$$P(X|Y) \approx \prod_{i=1}^{n}(P(x_i|f_j) + \epsilon_{j,i}), \tag{2}$$

where $j = 1, 2, ..., m$, $i = 1, 2, ..., n$, $P(x_i|f_j)$ is an univariate model that represents the inverse mapping from objective f_j to decision variable x_i, and $\epsilon_{j,i}$ is an error term. For convenience, it is assumed that $\epsilon_{j,i} \sim \mathcal{N}(0, (\sigma_n)^2)$ can be captured by additive Gaussian noise. Consequently, each univariate model together with the error term is realized using Gaussian process [13], which has the advantage of modeling both the global regularity and the local randomness in the distribution of the non-dominated solutions during the search. It is worth noting that although the decomposition strategy does not take into account the variable linkages explicitly, in our algorithm, a *random grouping method* has been adopted to implicitly learn the correlations between different decision variables by relating a number of decision variables with each objective. For example, for a three-objective MOP, if the group size is 2, three groups of models will be generated, each containing two univariate models. The reader is referred to [2].

In order to facilitate the inverse modelling, some pre-defined uniformly distributed reference vectors are used to partition the objective space into a number of subregions. To generate these reference vectors, the uniformly distributed points are firstly generated on a unit hyperplane and then mapped to a unit hypersphere, as shown in Fig. 1. With the pre-defined reference vectors (or subregions), the entire population can be partitioned into a number of subpopulations by associating the candidate solutions with different reference vectors. To associate each candidate solution with a particular reference vector, the angle (in the objective space) between each candidate solution and each reference vector is calculated, and a candidate solution is associated with a reference vector if

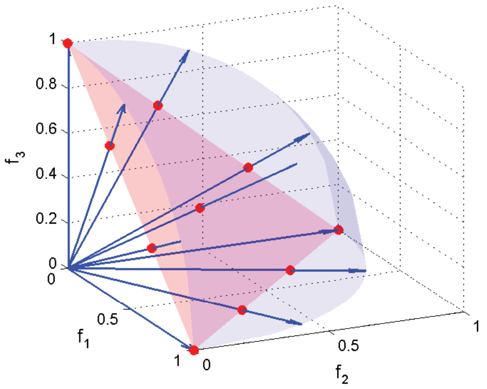

Fig. 1. An example of how to generate a number of 15 uniformly distributed reference vectors in a 3-objective space.

and only if the angle between the candidate solution and the reference vector is smallest among all reference vectors.

Partitioning a population using reference vectors in the objective space was first suggested in [10], which has also been adopted in a few other recently proposed algorithms such as NSGA-III [4,8]. However, the population partition strategy in NSGA-III is to use reference points distributed on a unit hyperplane in the objective space to guide the convergence of the population, and as a consequence, each individual in the population is expected to converge to a corresponding reference point. By contrast, our method is motivated to partition the actual objective space by setting a number of reference vectors, and around each reference vector, a subpopulation is maintained in the subregion defined by this reference vector. In each subregion, promising candidate solutions are selected using non-dominated sorting and crowd distance [5]. Inverse models are then built using the selected candidate solutions as the training data. Therefore, reproduction is operated in each subregion by sampling the inverse models built for this region. At the end of each generation, the offspring generated in each subregion is combined together to create the parent population for the next generation.

As shown in Fig. 2, the main operations of IM-MOEA, i.e., non-dominated sorting, selection, inverse modeling and reproduction, are all carried out within each subpopulation once the entire population is partitioned. Therefore, the reference vectors, which directly determine how the population is partitioned, play a central role in IM-MOEA.

3 Adaptive Reference Vector Generation

The adoption of uniformly distributed reference vectors in the original IM-MOEA is based on the implicit assumption that the Pareto front of the MOP

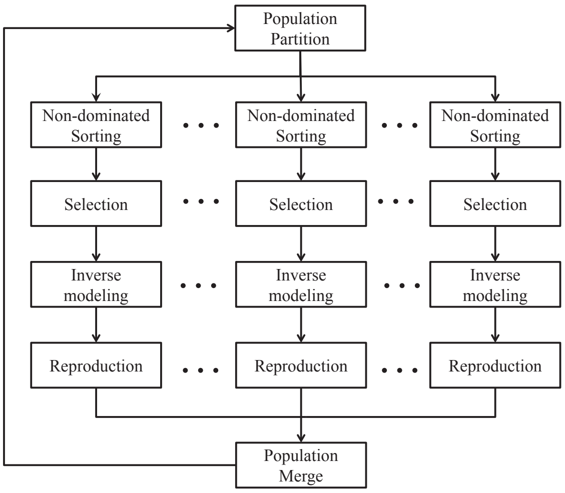

Fig. 2. The framework of the IM-MOEA.

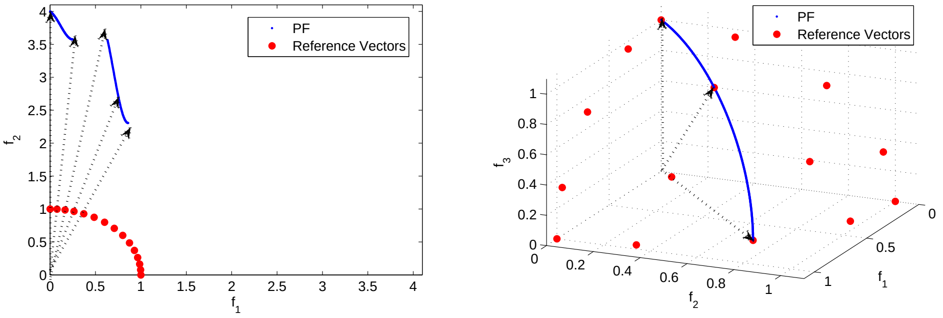

(a)

(b)

Fig. 3. Examples where there exist *invalid* reference vectors: (a) only 5 out of 15 reference vectors are covered by the Pareto front of 2-objective DTLZ7; (b) only 3 out of 15 reference vectors are covered by the Pareto front of 3-objective DTLZ5.

is uniformly distributed in the whole objective space. This assumption may be impractical for many real-world MOPs. In this work, without the loss of generality, we use two benchmark functions in the DTLZ test suite [6], namely, DTLZ5 and DTLZ7, as examples to examine the effectiveness of the proposed strategy for adaptively generating reference vectors.

DTLZ7 is a typical MOP with a disconnected Pareto front consisting of 2^{m-1} disconnected segments, where m is the objective number. For example, as shown in Fig. 3(a), the Pareto front of a 2-objective DTLZ7 consists of two Pareto optimal regions. Moreover, both segments of the Pareto front distribute in the left part of the objective space, resulting in a large region in the objective space that has no Pareto optimal solutions. In this case, for 15 reference vectors uniformly distributed in the objective space, only 5 out of 15 are associated with the Pareto optimal solutions, leaving 10 reference vectors unused.

Another typical MOP that suffers from such problems is DTLZ5. This MOP has a degenerate Pareto front, i.e., the Pareto front is always a curve regardless of the dimensionality of the objective space. As shown in Fig. 3(b), the Pareto front of the 3-objective DTLZ5 is a curve in the middle of the objective space. In this case, again, most of the uniformly distributed reference vectors are not in use (in this example only 3 out of 15 reference vectors), which will give rise to considerable waste of computational resources.

In practice, the distribution of the true Pareto front is usually not known beforehand. Therefore, in order to effectively use all reference vectors and the associated computational resources, it is essential to detect the distribution of the candidate solutions during the search and then adapt the distribution of the reference vectors accordingly. As mentioned before, since one candidate solution is associated with a reference vector if and only if their positions in terms of the angle between them in the objective space are closest, the density of the solutions in a subregion can be easily estimated by counting the number of candidate solutions associated with each reference vector. In this way, we are able to rank the reference vectors according to the numbers of candidate solutions associated with them. If one reference vector is associated with a small number of candidate solutions, it indicates that the density of candidate solutions in the subregion specified by this reference vector is small. In the extreme case, as shown in Fig. 3, no candidate solutions will be associated with a reference vector, which is termed an *invalid vector* in this work.

Based on the ranking of the reference vectors, we can also remove some undesirable reference vectors. Intuitively, the reference vectors associated with few candidate solutions should be removed. However, our empirical observations show that removing these reference vectors too early may cause a severe loss of population diversity, as in the early stage of search, exploration of the whole search landscape can be more important than exploitation, whereas in the late stage of the search, exploitation is more desirable. In order to adapt the reference vectors to the different preferences at different search stages, we divide the search into two phases: exploration phase and exploitation phase. In the exploration phase, since the reference vectors are expected to be widespread, the reference vectors associated with too many candidate solutions are preferentially removed. By contrast, in the exploitation phase, reference vectors associated few candidate solutions are removed since a high density of reference vectors in the subregions to be exploited can accelerate convergence.

Procedure 1 The pseudo code of the adaptive reference vector generation strategy.

Input:current generation t, max generation max_t, the current reference vectors $V(t)$, current population $P(t)$

Output:the adapted reference vectors $V(t+1)$

```
randomly generate a new reference vector $v_n$;
calculate the numbers of candidate solutions in $P(t)$ that associated with each reference vector in $V(t)$;
if $t < \theta * max\_t$ then
    /*exploration phase*/
    remove the reference vector in $V(t)$ which is associated with the *maximal* number of candidate solutions and replace it with $v_n$;
else
    /*exploitation phase*/
    remove the reference vector in $V(t)$ which is associated with the *minimal* number of candidate solutions and replace it with $v_n$;
end if
$V(t+1) = V(t)$;
```

To maintain a relatively stable distribution of the reference vectors, only one reference vector, i.e., the one is ranked first or last, will be removed in each generation. Meanwhile, a new reference vector is randomly generated to replace the removed one. In this way, in each generation, the extreme reference vector will be replaced with a new, randomly generated reference vector. The procedure of the adaptive reference vector generation strategy is summarized in Algorithm 1. It can be seen that a control parameter θ is introduced to determine at which generation the exploration stage is switched to exploitation. In Section 4, some preliminary empirical studies have been carried out to investigate the influence of parameter θ on the search performance.

4 Simulation Results

To assess the performance of the proposed adaptive reference vector generation strategy, IM-MOEA with the adaptive strategy, denoted as A-IM-MOEA hereafter, is compared with the original IM-MOEA on four three-objective DTLZ benchmark MOPs, including DTLZ1, DTLZ2, DTLZ5 and DTLZ7. The first two MOPs (DTLZ1, DTLZ2) have a uniformly distributed Pareto front, while the other two MOPs (DTLZ5, DTLZ7) have a degenerate and a disconnected Pareto front, respectively, refer to Fig. 3. The specific settings of these four MOPs follow the recommendations in [6].

In IM-MOEA and the proposed adaptive A-IM-MOEA, there are three parameters to be specified: the population size, the number of reference vectors, denoted as K_r hereafter, the group size (for random grouping). The population size is set to 150 in all the experiments. To investigate the sensitivity of the proposed adaptive strategy to K_r, different settings ($K_r = 10$, $K_r = 15$ and $K_r = 28$) have been used in the comparisons with IM-MOEA. In addition, it is worth noting that the setting of the group size is dependent on the number of decision variables. Since the numbers of decision variables of three-objective DTLZ1, DTLZ2, DTLZ5 and DTLZ7 (7, 11, 11 and 21, respectively) are small, the group size is simply set to 1.

The inverted generational distance (IGD) [1] is used as the performance indicator in the performance comparisons:

$$IGD(P^*, P) = \frac{\sum_{v \in P^*} d(v, P)}{|P^*|}, \tag{3}$$

where P^* is a set of uniformly distributed solutions along the true Pareto front, and P is an approximation, $d(v, P)$ is the minimum Euclidean distance from the point v to P. The IGD metric is able to measure both diversity and convergence of P if $|P^*|$ is large enough, and a smaller IGD value indicates a better performance. In this work, a number of 500 uniformly distributed points are selected for each benchmark MOP to be P^*.

Table 1. Statistical results of IGD values obtained by A-IM-MOEA and IM-MOEA (mean values in the first line and standard deviations in the second line). If one result is statistically significantly better than the other one, it is highlighted.

K_r	Algorithm	DTLZ1	DTLZ2	DTLZ5	DTLZ7
10	A-IM-MOEA	5.91E-02 1.37E-02	6.21E-02 3.83E-03	**6.05E-03** **6.24E-04**	**6.98E-02** **1.01E-02**
	IM-MOEA	4.97E-02 9.60E-03	**5.67E-02** **1.54E-03**	1.77E-02 1.94E-03	2.07E-01 2.30E-02
15	A-IM-MOEA	**3.83E-02** **6.19E-03**	5.44E-02 1.84E-03	**4.77E-03** **6.23E-04**	**6.00E-02** **7.07E-03**
	IM-MOEA	5.04E-02 1.71E-02	**5.06E-02** **7.54E-04**	1.50E-02 1.56E-03	1.18E-01 8.75E-03
28	A-IM-MOEA	**4.55E-02** **1.12E-02**	5.18E-02 2.40E-03	**4.04E-03** **5.53E-04**	**5.46E-02** **3.79E-03**
	IM-MOEA	4.76E-02 1.07E-02	4.88E-02 5.63E-04	1.06E-02 1.12E-03	8.96E-02 5.03E-03
+ / ≈ / −		2 / 1 / 0	0 / 2 / 1	3 / 0 / 0	3 / 0 / 0

The statistical results presented in this section are obtained from 20 independent runs. In each independent run, a maximum of 50000 fitness evaluations is used as a termination criterion for DTLZ2, DTLZ5 and DTLZ7. For DTLZ1, 150000 fitness evaluations are used. To compare the results obtained by A-IM-MOEA and IM-MOEA, the Wilcoxon rank sum test is adopted at a significance level of 0.05. As a result of the Wilcoxon rank sum test, "+" means that the IGD values obtained by A-IM-MOEA are statistically significantly smaller than those obtained by IM-MOEA; "−" means that the IGD values obtained by A-IM-MOEA are statistically significantly larger than those obtained by IM-MOEA; and "≈" means that there is no statistically significant difference between the IGD values obtained by A-IM-MOEA and IM-MOEA.

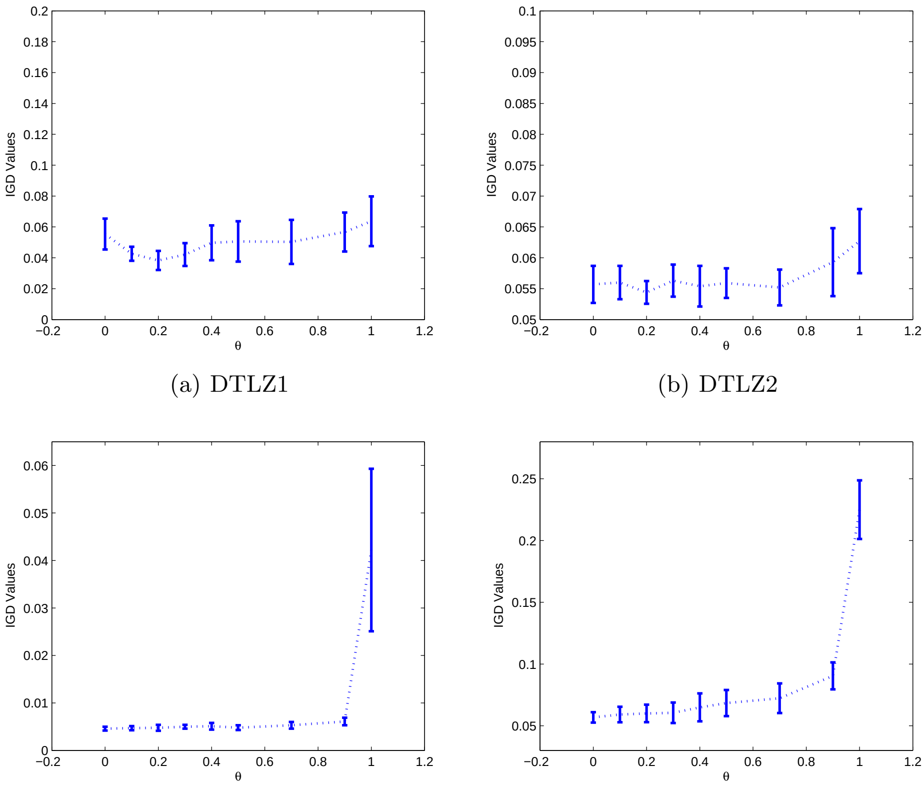

(a) DTLZ1

(b) DTLZ2

(c) DTLZ5

(d) DTLZ7

Fig. 4. The statistical results of the IGD values obtained by A-IM-MOEA with 15 reference vectors and different settings of θ. In this figure, error bars are used to present the mean and standard deviation.

Before comparing the performance of A-IM-MOEA and IM-MOEA, some investigations regarding the setting of θ will be conducted. *theta* determines the generation at which the exploration search stage is to be switched to exploitation, refer to Algorithm 1. As shown in Fig. 4, different settings of θ may have different impacts on different benchmark problems. For DTLZ5 and DTLZ7, which have a degenerate and a disconnected PF, respectively, it seems that the performance of

A-IM-MOEA is relatively insensitive to the settings of θ, as long as it is not larger than 0.7. This might be due to the fact that DTLZ5 and DTLZ7 are uni-modal, exploration has no significant effect on the search performance. Exploitation, which mainly contributes to convergence, can be important as the Pareto fronts of these two MOPs are not uniformly distributed in the objective space. By contrast, since the fitness landscape of DTLZ1 contains a large number of local optima, sufficient exploration becomes more important. It can be seen in Fig. 4 (a) that when θ is between 0.1 and 0.3, the standard deviation of IGD is smaller compared to that in other settings, which implies a more stable performance of A-IM-MOEA. Among the four benchmark MOPs, DTLZ2 is uni-modal and has a uniformly distribute Pareto front. Therefore, the performance of A-IM-MOEA is not very sensitive to the settings of θ either.

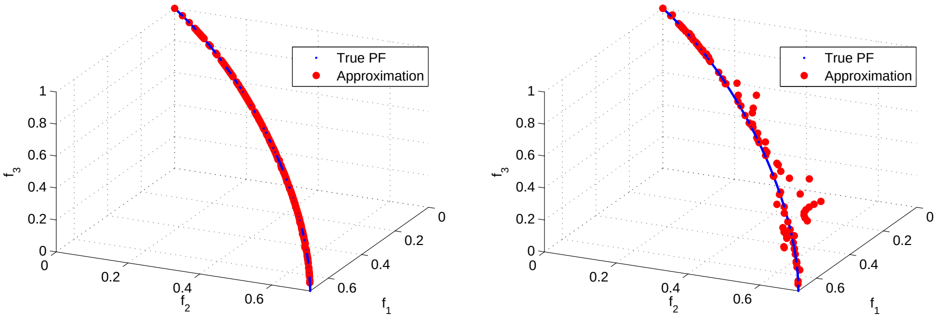

(a) A-IM-MOEA

(b) IM-MOEA

Fig. 5. The solutions obtained by A-IM-MOEA and IM-MOEA with 28 reference vectors on DTLZ5 in the final population of the best singe run.

Based on the empirical investigations on the setting of θ, we use $\theta = 0.2$ for all experiments for comparing A-IM-MOEA and IM-MOEA. The statistical results obtained by A-IM-MOEA and IM-MOEA are summarized in Table 1. It can be seen that A-IM-MOEA significantly outperforms IM-MOEA on DTLZ5 and DTLZ7, regardless of the number of reference vectors. As evident from Fig. 5, the solutions obtained by A-IM-MOEA show significantly better convergence. It is because the reference vectors in A-IM-MOEA have been successfully adapted, thus increasing the sampling density around the true Pareto front rather than the entire objective space. To verify this statement, the reference vectors are plotted together with the true Pareto front. To better visualize the adapted distribution of the reference vectors, the points are mapping into a 2-D (f_1 and f_2) plane, as shown in Fig. 6. It can be seen that the reference vectors in A-IM-MOEA are mostly distribute around the true Pareto front, whilst the reference vectors in IM-MOEA, without any adaption, still uniformly distributed in the entire objective space.

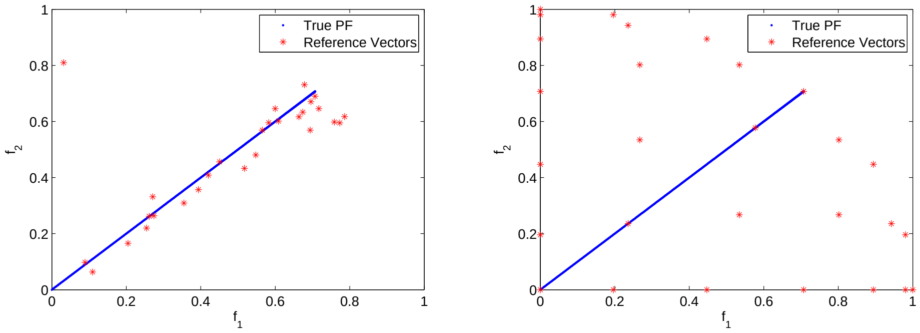

(a) A-IM-MOEA (b) IM-MOEA

Fig. 6. The solutions obtained by A-IM-MOEA and IM-MOEA with 28 reference vectors on DTLZ5 in the final population of the best singe run. To ease the observations, the points are mapping into a 2-D (f_1 and f_2) plane from the 3-D objective space.

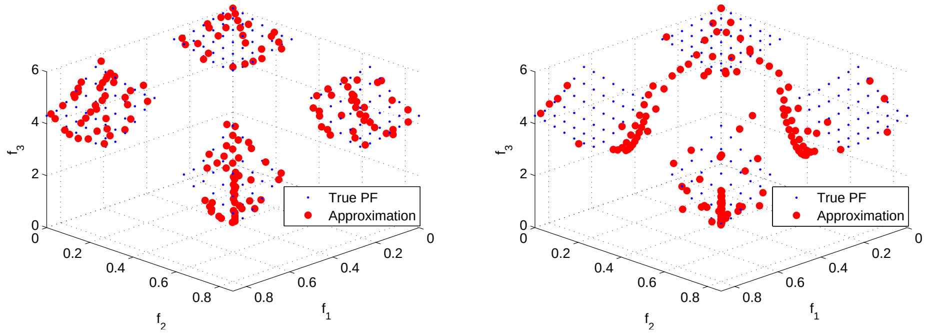

(a) A-IM-MOEA (b) IM-MOEA

Fig. 7. The solutions obtained by A-IM-MOEA and IM-MOEA with 28 reference vectors on DTLZ7 in the final population of the best single run.

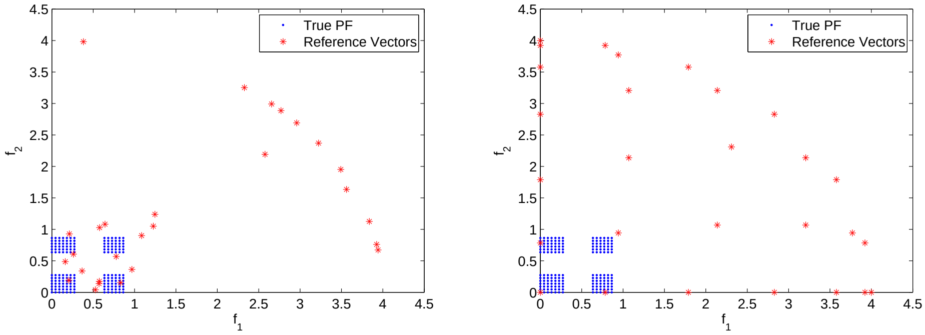

(a) A-IM-MOEA

(b) IM-MOEA

Fig. 8. The solutions obtained by A-IM-MOEA and IM-MOEA with 28 reference vectors on DTLZ7 in the final population of the best singe run. The points are mapping into a 2-D (f_1 and f_2) plane from the 3-D objective space for better visualization.

Similar observations can be made from the results on DTLZ7 as well. As evident from Fig. 7, the solutions obtained by A-IM-MOEA show a promising distribution, while most of the solutions obtained by IM-MOEA distribute on the edges of the true Pareto front consisting of four disconnected piecewise segments. In addition to disconnection, the Pareto front of DTLZ7 shows significant bias on the m-th objective, thus resulting the distribution of the Pareto front centralized close to the third axis (f_3) in a 3-D objective space. These properties raise considerable difficulties for IM-MOEA which adopts a uniformly distributed reference vectors. By contrast, the adaptive reference vector generation strategy adopted in A-IM-MOEA has significantly better efficiency, as indicated in Fig. 8.

From the statistical results in Table 1, another interesting observation is the comparable performance of A-IM-MOEA and IM-MOEA on the other two MOPs, DTLZ1 and DTLZ2, which have a uniformly distributed Pareto front. On DTLZ1, the performance of A-IM-MOEA is slightly better while on DLTZ2, IM-MOEA shows slightly better performance. We surmise that A-IM-MOEA is outperformed by IM-MOEA on DTLZ2 because the search process is more or less disturbed by the adaptively changing reference vectors. By contrast, the predefined uniformly distributed reference vectors adopted in IM-MOEA can lead to a more stable search. For MOPs with a uniformly distributed Pareto front, where a predefined set of reference vectors is more desirable, IM-MOEA can outperform A-IM-MOEA. Nevertheless, for DTLZ1, which has a uniformly distributed Pareto front as well, this observation does not hold. This might be attributed to the fact that DTLZ1 is a multi-modal MOP, which may require a higher degree of population diversity for a better exploration. In this case, A-IM-MOEA can be more promising as the adaptively changing reference vectors can generate higher population diversity than the predefined reference vectors. As shown in Fig. 9, the adaptively changing reference vectors have enhanced the

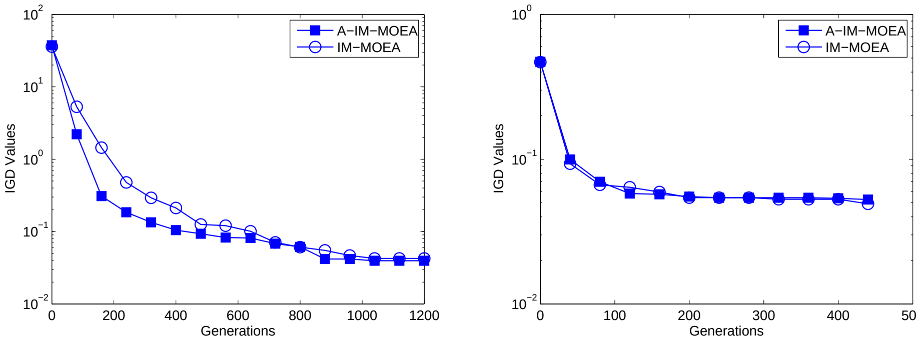

(a) DTLZ1

(b) DTLZ2

Fig. 9. The convergence profiles of the IGD values in the best single run with 15 reference vectors of A-IM-MOEA and IM-MOEA respectively.

convergence speed of A-IM-MOEA in the exploration stage on DTLZ1. However, on DTLZ2, which is a uni-modal MOP, the convergence profiles of A-IM-MOEA and IM-MOEA show little difference.

5 Conclusion

An adaptive reference vector generation strategy is proposed in this paper, which has shown to be promising on two MOPs having a discrete or non-uniform Pareto front. In addition, the MOEA using the proposed strategy performs comparably with the one using uniformly distributed reference vectors on MOPs having a uniform Pareto front distributed in the whole objective space. In addition, an interesting observation is that the adaptive reference vectors are able to generate better population diversity to enhance the performance of IM-MOEA on multi-modal MOPs like DTLZ1.

In the future, the performance of the proposed adaptive strategy for generating reference vectors will be further assessed on additional MOPs. For example, it can be interesting to see how it performs on constrained MOPs, where the Pareto front is irregular as well. The mechanism for switching between exploration and exploitation also needs further examination.

Acknowledgment

This work was supported in part by the Honda Research Institute Europe GmbH and the Joint Research Fund for Overseas Chinese, Hong Kong and Macao Scholars of the National Natural Science Foundation of China (Grant No. 61428302).

References

1. Bosman, P.A., Thierens, D.: The balance between proximity and diversity in multiobjective evolutionary algorithms. IEEE Transactions on Evolutionary Computation 7(2), 174–188 (2003)
2. Cheng, R., Jin, Y., Narukawa, K., Sendhoff, B.: A multiobjective evoltuionary algorithm using Gaussian process based inverse modeling. IEEE Transactions on Evolutionary Computation (submitted in 2014)
3. Cornell, J.A.: Experiments with mixtures: designs, models, and the analysis of mixture data. John Wiley & Sons (2011)
4. Deb, K., Jain, H.: An evolutionary many-objective optimization algorithm using reference-point based non-dominated sorting approach, part I: Solving problems with box constraints. IEEE Transactions on Evolutionary Computation 18(4), 577–601 (2013)
5. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation 6(2), 182–197 (2002)
6. Deb, K., Thiele, L., Laumanns, M., Zitzler, E.: Scalable multi-objective optimization test problems. In: Proceedings of the Congress on Evolutionary Computation. pp. 825–830. IEEE (2002)
7. Giagkiozis, I., Fleming, P.J.: Pareto front estimation for decision making. Evolutionary Computation (accepted in 2014)
8. Jain, H., Deb, K.: An evolutionary many-objective optimization algorithm using reference-point based non-dominated sorting approach, part II: Handling constraints and extending to an adaptive approach. IEEE Transactions on Evolutionary Computation 18(4), 602–622 (2013)
9. Jin, Y., Sendhoff, B.: Connectedness, regularity and the success of local search in evolutionary multi-objective optimization. In: Proceedings of the IEEE Congress on Evolutionary Computation. vol. 3, pp. 1910–1917. IEEE (2003)
10. Liu, H.L., Gu, F., Zhang, Q.: Decomposition of a multiobjective optimization problem into a number of simple multiobjective subproblems. IEEE Transactions on Evolutionary Computation 18(3), 2450–455 (2014)
11. Martí, L., Garcia, J., Berlanga, A., Coello Coello, C., Molina, J.M.: On current model-building methods for multi-objective estimation of distribution algorithms: Shortcommings and directions for improvement. Department of Informatics, Universidad Carlos III de Madrid, Madrid, Spain, Technical Report GIAA2010E001 (2010)
12. Okabe, T., Jin, Y., Sendoff, B., Olhofer, M.: Voronoi-based estimation of distribution algorithm for multi-objective optimization. In: Proceedings of the IEEE Congress on Evolutionary Computation. vol. 2, pp. 1594–1601. IEEE (2004)
13. Rasmussen, C.E.: Gaussian processes for machine learning. MIT Press (2006)
14. Zhang, Q., Li, H.: MOEA/D: A multiobjective evolutionary algorithm based on decomposition. IEEE Transactions on Evolutionary Computation 11(6), 712–731 (2007)
15. Zhang, Q., Zhou, A., Jin, Y.: RM-MEDA: A regularity model-based multiobjective estimation of distribution algorithm. IEEE Transactions on Evolutionary Computation 12(1), 41–63 (2008)
16. Zhou, A., Qu, B.Y., Li, H., Zhao, S.Z., Suganthan, P.N., Zhang, Q.: Multiobjective evolutionary algorithms: A survey of the state of the art. Swarm and Evolutionary Computation 1(1), 32–49 (2011)