The Domination Game: Dilating Bubbles to Fill Up Pareto Fronts

Vasco Coelho
Department of Informatics, Systems and Communication
University of Milano-Bicocca, Milan, Italy
v.coelho@campus.unimib.it

Daniele M. Papetti
Department of Informatics, Systems and Communication
University of Milano-Bicocca, Milan, Italy
daniele.papetti@unimib.it

Andrea Tangherloni
Bocconi Institute for Data Science and Analytics
Department of Computing Sciences
Bocconi University, Milan, Italy
andrea.tangherloni@unibocconi.it

Paolo Cazzaniga
Department of Human and Social Sciences
University of Bergamo, Bergamo, Italy
paolo.cazzaniga@unibg.it

Daniela Besozzi
Department of Informatics, Systems and Communication
University of Milano-Bicocca, Milan, Italy
daniela.besozzi@unimib.it

Marco S. Nobile
Department of Environmental Sciences, Informatics and Statistics
Ca' Foscari University of Venice, Venice, Italy
marco.nobile@unive.it

2023 IEEE Congress on Evolutionary Computation (CEC) | 979-8-3503-1458-8/23/$31.00 ©2023 IEEE | DOI: 10.1109/CEC53210.2023.10253982

***Abstract*—Multi-objective optimization algorithms might struggle in finding optimal dominating solutions, especially in real-case scenarios where problems are generally characterized by non-separability, non-differentiability, and multi-modality issues. An effective strategy that already showed to improve the outcome of optimization algorithms consists in manipulating the search space, in order to explore its most promising areas. In this work, starting from a Pareto front identified by an optimization strategy, we exploit Local Bubble Dilation Functions (LBDFs) to manipulate a locally bounded region of the search space containing non-dominated solutions. We tested our approach on the benchmark functions included in the DTLZ and WFG suites, showing that the Pareto front obtained after the application of LBDFs is most of the time characterized by an increased hyper-volume value. Our results confirm that LBDFs are an effective means to identify additional non-dominated solutions that can improve the quality of the Pareto front.**

***Index Terms*—Multi-objective Optimization, Global Optimization, Pareto Front, Search Space Manipulation, Local Bubble Dilation Functions**

I. Introduction

The last decade has witnessed an increasing number of studies regarding real-world optimization problems [1], [2], ranging from civil engineering to manufacturing and intelligent system design [3], [4], hydrology [5], network analysis [6], medical imaging [7], and life sciences [8]–[12]. Most of these problems can be modeled using multiple conflicting objectives, making Multi-Objective Optimization (MOO) algorithms fundamental for identifying optimal solutions [13]. In this context, the family of Multi-Objective Evolutionary Algorithms (MOEAs) is the most used. Being evolutionary approaches, MOEAs evolve a population of randomly created individuals to approximate the Pareto optimal front [14]. An improvement in one of the objectives generally worsens some of the others; thus, most of the time, a solution that simultaneously optimizes all the objectives does not exist. However, there is a set of Pareto optimal solutions (i.e., the optimal Pareto set) that contains the best trade-offs in terms of objective values [14].

Among the plethora of MOEAs, it is worth mentioning the Strength Pareto Evolutionary Algorithm (SPEA) [15] and the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [16]. SPEA exploits a regular population and an external set (i.e., an archive of non-dominated solutions) that are iteratively updated. The individuals undergo recombination and mutation operators to generate offspring. In addition, a clustering technique is applied to preserve the characteristics of the non-dominated solutions when the archive size exceeds a predefined limit. SPEA2 is an improved version of SPEA that integrates a fine-grained fitness assignment strategy, a density estimation technique, and an enhanced archive truncation method [17]. NSGA-II is based on the general scheme of Genetic Algorithms but exploits a different selection strategy, whereby the individuals are selected based on both the Pareto front and a crowding distance (i.e., the Manhattan distance in the objective space) in the generated splitting front. NSGA-III is an improved version of NSGA-II that integrates a selection from the splitting front based on a set of reference directions on a unit simplex [18], [19].

MOEAs often show difficulties in solving real-world optimization problems that exhibit pathological fitness landscape characteristics, such as non-separability, non-differentiability, and multi-modality [20]. For this reason, several research studies tried to improve their performance by adding new features or generating surrogate models of the fitness function [21]–[24]. Surrogate models allow for evaluating simpler fitness landscapes; however, their application is not always cost-effective as their construction generally represents a difficult task. Other research studies employ diversity-preserving mechanisms to construct uniformly spread Pareto front approximations. In [25], for instance, the authors propose an archiving strategy to improve the discretization of the Pareto front geometry by optimizing some physical interactions (e.g., Coulomb's law) between pairs of candidate solutions. This strategy can be coupled with any multi-objective evolutionary algorithm to keep track of candidate solutions discarded by the selection operator.

Another possible approach is the direct application of transformation or modification strategies to the search space, to the aim of improving the quality of the sampling procedure of candidate solutions by avoiding non-promising fitness landscape regions. For instance, the Shrinking Space Technique [26] and Space Transformation Search [27] utilize shrinking and transformation, respectively, during the optimization phase to focus on promising areas [26], [28], [29]. A different strategy relying on Dilation Functions (DFs) has been recently introduced to re-map the original search space onto a dilated search space [30]. The proposed DFs are used to expand the most promising regions and compress the regions characterized by poor fitness values. Considering that DFs are problem-dependent and require prior knowledge of the fitness landscape, an automatic method has been proposed to evolve optimal DFs [31]. An extension of the DFs was then proposed in [32], where the authors introduced the Local Bubble Dilation Functions (LBDFs), which apply a transformation to the fitness landscape in a locally bounded region, rather than dilating the entire search space. In particular, LBDFs perform mappings of solutions of the search space into other solutions of the search space, and they can be applied to both single-objective and MOO problems.

In the case of single-objective problems, the application of an LBDF is straightforward: the fitness value of the dilated solution is the fitness value of the re-mapped point in the original landscape. As a consequence, the landscape directly reflects the applied manipulation by showing different fitness values in the dilated region. When applying LBDFs to MOO problems, since the manipulation is performed at the candidate solution level (i.e., the search space), the effects are simultaneously reflected in all landscapes related to the different objectives. Figure 1 shows an LBDF applied to the DTLZ1 benchmark problem to dilate both landscapes of its two objectives. In the context of MOO problems, the purpose of applying LBDFs slightly differs from the purpose of applying them on single-objective problems. In fact, while in single-objective optimization LBDFs aim at dilating the basin of attraction where the global optimum lies, in multi-objective problems the aim is to dilate regions containing non-dominated solutions, rather than regions that improve the quality of only some of the landscapes.

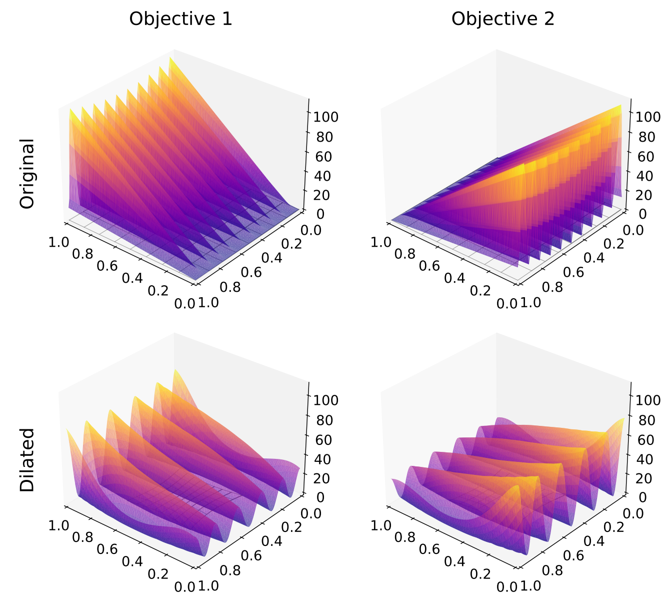

Fig. 1. Example of an LBDF applied to locally manipulate the search space of the two objectives of the DTLZ1 benchmark problem.

In this work, we exploit LBDFs to dilate the search space of MOO problems to improve the quality of the solutions in the Pareto set. In particular, given a Pareto set, LBDFs are applied to specific regions of the search space, marked as promising areas by analyzing the location of the already identified non-dominated solutions. To assess the effectiveness of our approach, we applied LBDFs to improve the quality—evaluated in terms of hyper-volume—of sets of non-dominated solutions of benchmark functions belonging to the DTLZ [33] and WFG [34] suites. All the analyzed problems are defined with both 2 and 3 objectives over 30 dimensions.

The paper is structured as follows. In Section II we remind the concepts of MOO, DFs, and LBDFs, and then we introduce the approach to fill a Pareto front by properly positioning LBDFs and identifying new non-dominated solutions. Section III shows the results of the application of LBDFs to the DTLZ and WFG benchmark suites. Our conclusive remarks are given in Section IV.

II. Methods

A. Multi-Objective Optimization

The main aim of a MOO algorithm is to determine a set of K non-dominated candidate solutions $\mathcal{P}^* = \{\boldsymbol{p}_1^*, \ldots, \boldsymbol{p}_K^*\}$, which cannot improve a single objective without affecting the others. Without loss of generality, let $f_1, \ldots, f_\Omega$ be Ω distinct objective functions to be simultaneously minimized in a D dimensional search space $S \subseteq \mathbb{R}^D$. A candidate solution $\boldsymbol{x_1} \in S$ is said *dominated* by a solution $\boldsymbol{x_2} \in S$ if:

- $f_i(\boldsymbol{x_2}) \leq f_i(\boldsymbol{x_1})$ for all $i = 1, \ldots, \Omega$;
- $f_j(\boldsymbol{x_2}) < f_j(\boldsymbol{x_1})$ for at least one index $j = 1, \ldots, \Omega$.

The non-dominance relation produces a partially ordered set of non-dominated solutions that can be used to identify the different Pareto sets, each containing all the solutions with the same ranking; the Pareto set containing the non-dominated solutions is usually referred to as the "first" Pareto set. The fitness values of the candidate solutions of a Pareto set can be represented in an objective space, which is a Ω-dimensional space whose i-th dimension represents the fitness values of f_i.

B. Basis Functions

The family of Basis Functions (BFs) includes all the bijective and monotonically increasing functions that map the unit interval into the unit interval. BFs can be composed to define Dilation Functions (DFs), which can alter the semantics of candidate solutions. In this work, we focus on the following BF [31], [32]: $h_\gamma(\boldsymbol{x}) = \boldsymbol{x}^\gamma$, with $\gamma \in (1, \infty)$, whose inverse function is defined as $h_\gamma^{-1}(\boldsymbol{x}) = h_{1/\gamma}(\boldsymbol{x})$. The parameter γ influences the intensity of the dilation effect: the further the curve is from the identity function, the stronger the dilation of the search space will be, as shown in Figure 2.

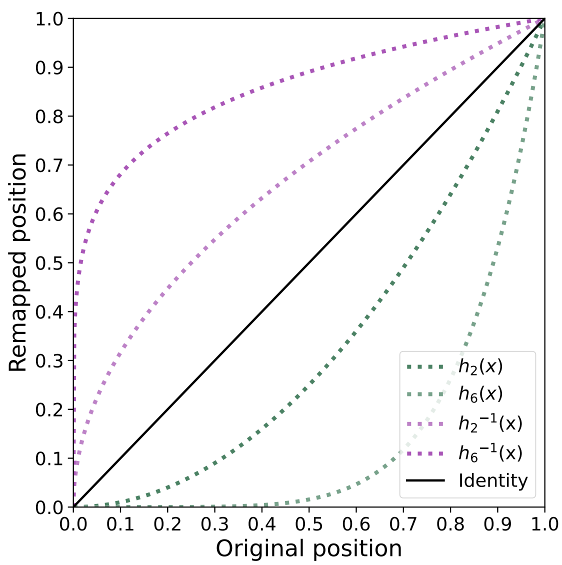

Fig. 2. Examples of Basis Functions h_γ and h_γ^{-1} with different values of γ.

C. Local Bubble Dilation Functions

Local Bubble Dilation Functions (LBDFs) perform local manipulations of the problem landscape by applying a BF inside a hyper-sphere [32]. The hyper-sphere is identified by a point $\boldsymbol{c}$ of the search space and a radius $r \in \mathbb{R}^+$. Formally, an LBDF $B_{h_\gamma,r,\boldsymbol{c}} : \mathbb{R}^D \to \mathbb{R}^D$ is defined as follows:

$$B_{h_\gamma,r,\boldsymbol{c}}(\boldsymbol{x}) = h_\gamma\left(\frac{||\boldsymbol{x}-\boldsymbol{c}||_2}{r}\right) \cdot \frac{\boldsymbol{x}-\boldsymbol{c}}{||\boldsymbol{x}-\boldsymbol{c}||_2} \cdot r + \boldsymbol{c}, \qquad (1)$$

for every $\boldsymbol{x} \in \mathbb{R}^D$ such that $||\boldsymbol{x}-\boldsymbol{c}||_2 < r$, and $B_{h_\gamma,r,\boldsymbol{c}}(\boldsymbol{x}) = \boldsymbol{x}$ otherwise.

If an LBDF leverages a concave up BF, such as $h_\gamma(\boldsymbol{x})$, an expansion is performed: the candidate solutions of the search space are mapped toward the center of the hyper-sphere, leading to an enhancement of the search of this area. *Expanding* LBDFs can be used to facilitate the exploration of promising regions of the search space. Conversely, if a concave down BF is used, such as $h_\gamma^{-1}(\boldsymbol{x})$, a compression is performed: the candidate solutions inside the hyper-sphere are mapped further from the center, thus reducing the explorability of the area. *Compressing* LBDFs can be used to limit the search of non-promising regions of the search space.

Since LBDFs perform local manipulations of the search space, it is possible to simultaneously apply a composition of multiple LBDFs. In principle, this might yield landscapes where multiple promising regions are expanded and many poor areas are compressed. In this work, we focus on the application of multiple *expanding* LBDFs to improve the quality of a set of non-dominated solutions.

D. Dilation Functions to Fill the Pareto Front

Effective manipulations of the search space require detailed knowledge of the characteristics of the fitness landscape, which is generally not available. For this reason, we propose a novel approach to improve the quality of an already computed Pareto set and potentially increase the diversity among the non-dominated solutions.

Given the first Pareto front $\mathcal{P}$ computed on N candidate solutions, for each pair of adjacent solutions $\boldsymbol{p}, \boldsymbol{q} \in \mathcal{P}$ in the objective space, an LBDF is applied between them in the search space. The center of this LBDF corresponds to the centroid between $\boldsymbol{p}$ and $\boldsymbol{q}$, while its radius is equal to half of the Euclidean distance between $\boldsymbol{p}$ and $\boldsymbol{q}$. Considering a two-objective problem, the concept of adjacency is defined by the ordering along one objective. For three or more objectives it is not possible to operate in the same way, so the adjacency is defined in terms of minimum Euclidean distance. After the positioning of LBDFs, M new candidate solutions are sampled from the sub-regions of the search space identified by the LBDFs. These M solutions, together with the candidate solutions in $\mathcal{P}$, are eventually considered to compute the new Pareto set, following the first Pareto set selection strategy described in [16]. Figure 3 shows the workflow of the approach proposed in this work.

The key advantage of this approach is that it determines the promising regions of the search space by analyzing the location of non-dominated solutions, without performing additional fitness evaluations. Moreover, by using the new information extracted from the dilated search space, the quality of the Pareto front can either improve or, at worst, remain the same. As a matter of fact, if the newly sampled solutions are all dominated, the Pareto set remains the same and corresponds to an identical hyper-volume. On the contrary, if one or more sampled solutions are non-dominated, then the value of the hyper-volume can only increase, thanks to the improvement of the Pareto front. The approach proposed in this paper is based on the idea that a non-dominated solution might lie in the search space region bounded by two solutions belonging to the first Pareto front. We expect that the expansion of the regions between two Pareto-optimal solutions will make the identification of new dominating solutions easier.

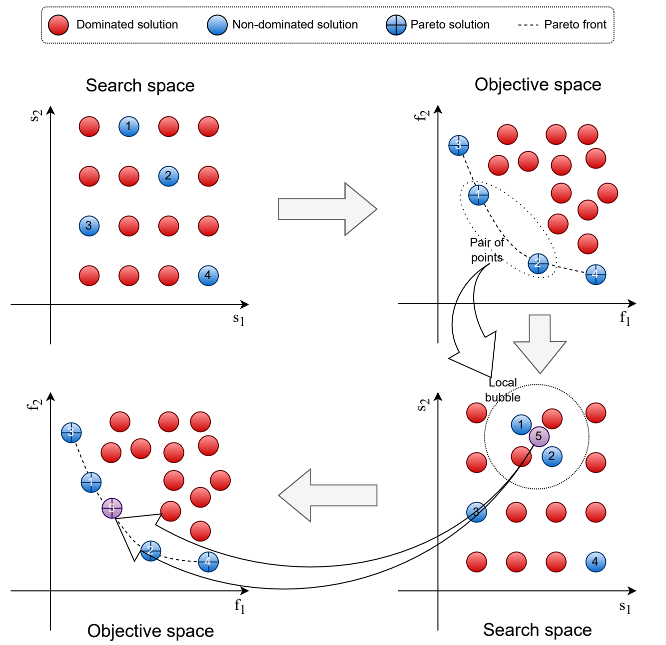

Fig. 3. Workflow of the proposed approach. Given a set of N candidate solutions in the original search space (top left), the first Pareto front is computed (the non-dominated solutions are denoted as blue dots). For each pair of adjacent solutions (e.g., solutions 1 and 2) in the objective space (top right), an LBDF is applied to the search space (bottom right). Additional candidate solutions are obtained from the dilated search space (e.g., solution 5) and the new Pareto set is computed (bottom left).

III. Results

We tested our approach on two common benchmark suites for MOO—namely, DTLZ and WFG—implemented in the `pymoo 0.6.0.1` Python library. For each benchmark function, we considered $D = 30$ dimensions and a number of objectives $\Omega \in \{2, 3\}$. The source code is available on GitHub: `https://github.com/Vsc0/the-domination-game`.

A. Two-Objectives Problems

For each benchmark function with $\Omega = 2$, we performed three different tests:

- `PF-100`: we generated 100 random candidate solutions, and calculated the corresponding Pareto front;
- `PF-50`: we selected 50 candidate solutions from the `PF-100` test, and calculated the corresponding Pareto front;
- `PF-LBDF`: we generated 50 new random candidate solutions in the search space regions manipulated by LBDFs, which are placed using the approach described in Section II-D. We then merged this population with the individuals selected in test `PF-50`, and calculated the resulting Pareto front.

Tests `PF-100` and `PF-LBDF` exploit the same budget in terms of candidate solutions evaluations, and the LBDFs are all based on the same dilation function h_γ, with $\gamma = 2$.

To quantify the impact of our strategy on the Pareto front, at the end of each test we calculated its hyper-volume as described in [35]. To compare different runs of the same benchmark function, we normalized the objective values in $[0, 1]$ by dividing each fitness value by the largest value observed for that objective. To this aim, we used $\mathbf{x} = (1, 1) \in \mathbb{R}^2$ as reference point of the hyper-volume. We performed 30 runs for each test to collect statistically significant results.

Figure 4 shows the box-plot representation of the distributions of the hyper-volume values. We applied the Mann–Whitney U test with Bonferroni correction [36]–[38] to evaluate any possible statistical difference among them. According to these results, the approach based on LBDFs is better than `PF-100` as the difference in the hyper-volume value is statistically significant in 13 cases out of 16. In the other cases, the approaches resulted to be statistically equivalent.

To visually clarify the impact of applying an LBDF, we show in Figure 5-A the Pareto fronts calculated using the results of tests `PF-100`, `PF-50`, and `PF-LBDF` on the DTLZ2 benchmark function, where the `PF-LBDF` outperformed the `PF-100`. If we compare the Pareto fronts obtained with tests `PF-100` (orange circles) and `PF-50` (blue squares), we can observe that their quality is basically equivalent. The application of the LDBFs has a beneficial effect since it allows for filling the gaps in the Pareto front described by the blue squares (e.g., between solutions 2–3, 4–5, and 7–8); the resulting approximation of the Pareto front is strongly improved and characterized by a higher hyper-volume.

A similar result is shown in Figure 5-C for the benchmark function WFG4. In this case, the application of the LBDFs between pairs of solutions found in the `PF-50` test led to the identification of a large number of additional Pareto-optimal solutions (blue crosses). Again, the `PF-LBDF` strategy led to a better characterization of the Pareto front, without using any actual optimization algorithm.

Figure 5-B shows the result obtained on the WFG1 benchmark function, which represents one of the cases where no statistically significant difference was observed. The first Pareto set includes a solution (i.e., 0) that is distant from all other solutions, which are instead clustered (those ranging from 1 to 8). What we expected, in this scenario, was to find additional non-dominated solutions between the two distant solutions (i.e., 0 and 1) and only a few non-dominated solutions close to the cluster in the Pareto set. On the contrary, the LBDFs revealed only one new non-dominated solution between 0 and 1 and ten solutions in the clustered region.

B. Three-Objectives Problems

For each benchmark function defined with $\Omega = 3$, we performed three different tests:

- `NSGAII-5000`: we executed NSGA-II with 50 individuals for 100 generations using a total budget of 5000 evaluations of the candidate solutions;

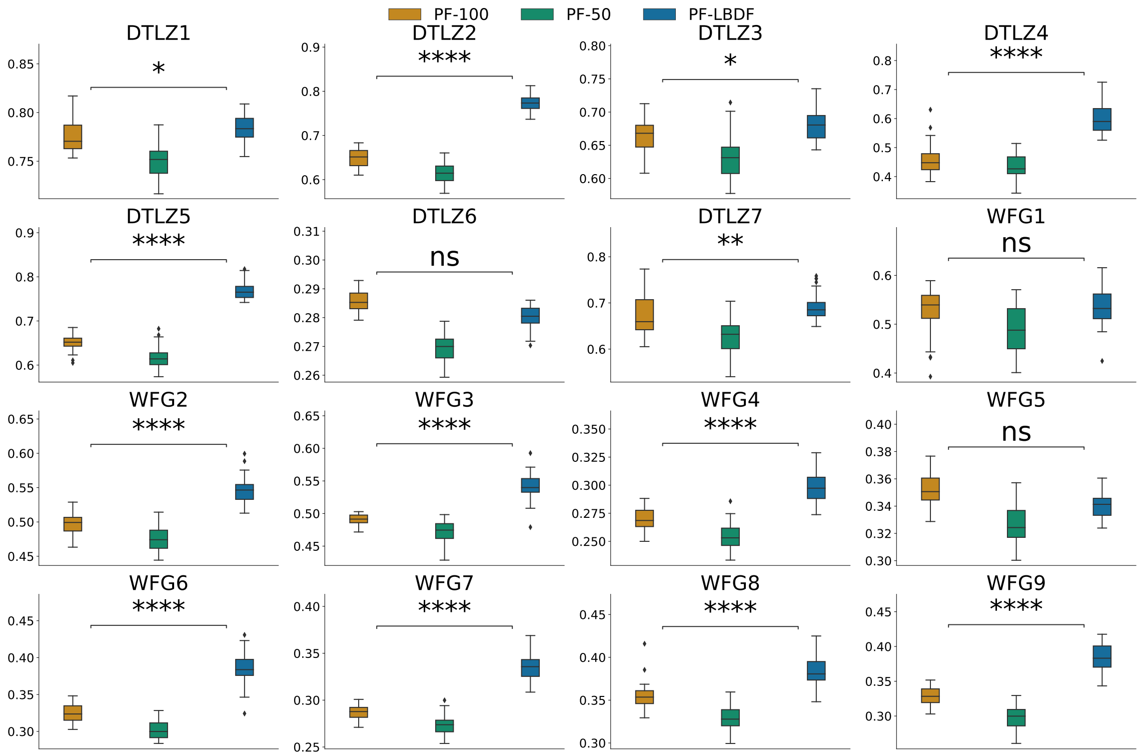

Fig. 4. The box-plots show the hyper-volume distribution across 30 runs of the proposed method on the DTLZ and WFG suites with 2 objectives. The Mann–Whitney U test with Bonferroni correction was applied to verify any statistically significant difference between the distributions built starting from the 100 randomly selected points either using the dilation (`PF-LBDF`) or not (`PF-100`) (**** p-value ≤ 0.0001, *** p-value ≤ 0.001, ** p-value ≤ 0.01, * p-value ≤ 0.05, ns p-value ≤ 1).

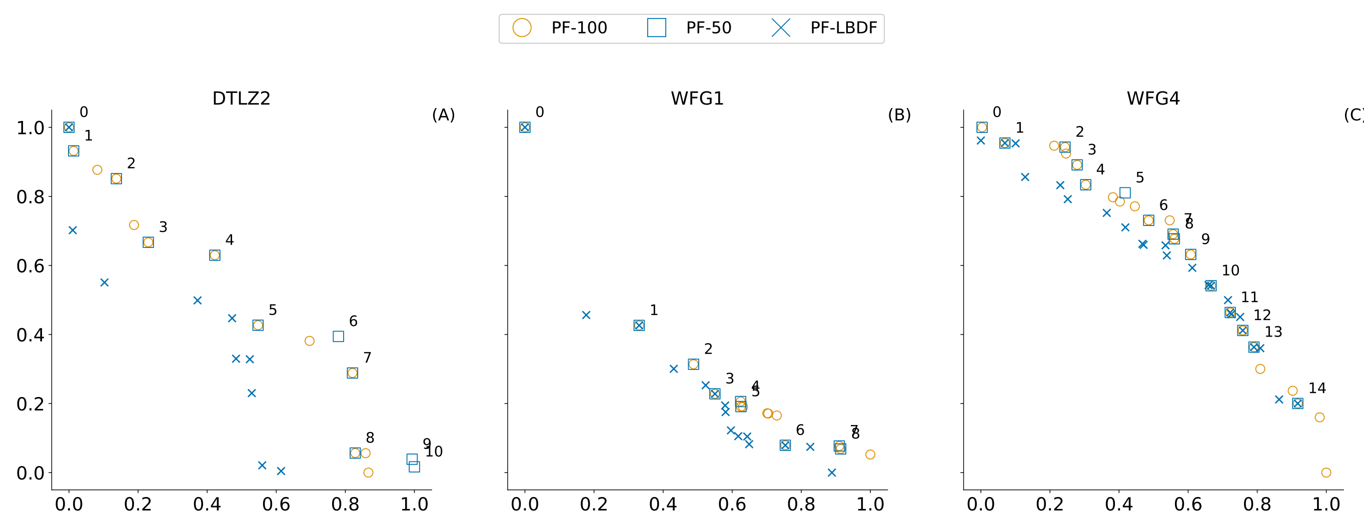

Fig. 5. Comparison of the Pareto fronts obtained on the DTLZ2, WFG4, and WFG1 benchmark problems. `PF-100`: the orange circles denote the non-dominated solutions obtained by considering 100 solutions randomly sampled in the original search space. `PF-50`: the blue squares represent the non-dominated solutions obtained by considering only 50 solutions of `PF-100`. `PF-LBDF`: the blue crosses indicate the non-dominated solutions obtained by considering 50 solutions randomly sampled in the dilated search space, in addition to the solutions of `PF-50`.

- `NSGAII-4750`: we executed NSGA-II with 50 individuals for 95 generations using a total budget of 4750 evaluations of candidate solutions;
- `LBDF-250`: we generated 250 new random candidate solutions in the search space regions, altered by means of LBDFs using the proposed approach. We then merged this population with the individuals created in test `NSGAII-4750`, and calculated the resulting Pareto front.

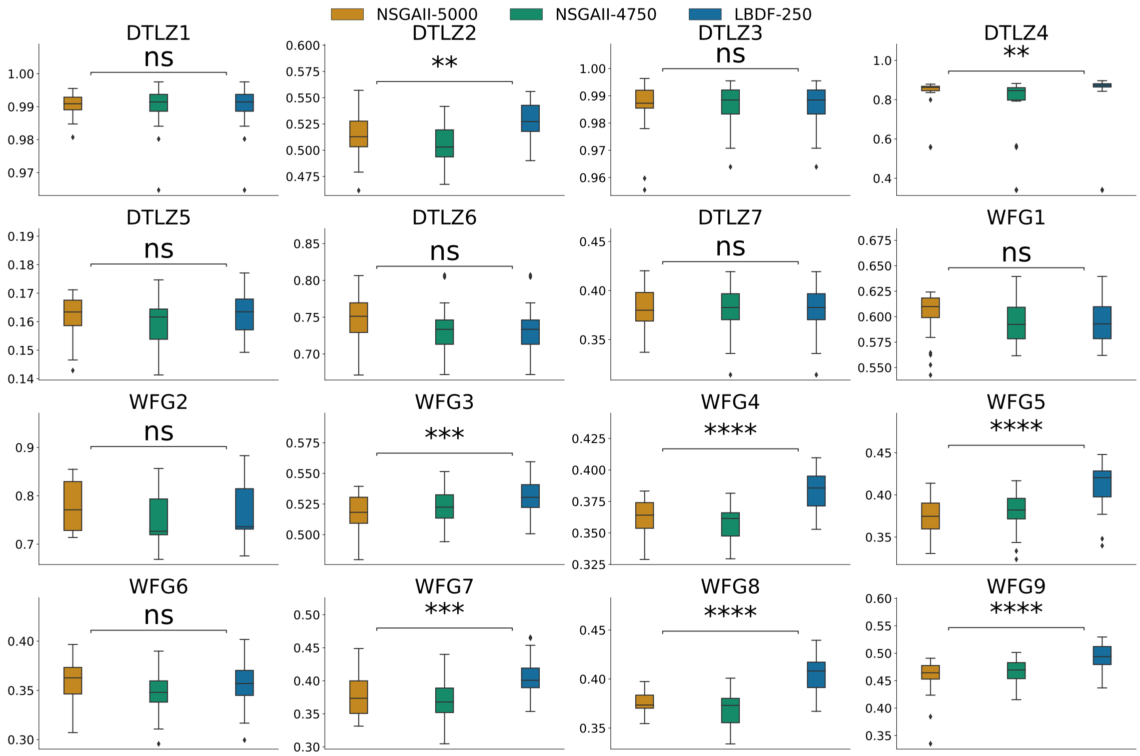

Fig. 6. The box-plots show the hyper-volume distribution across 30 executions of `NSGAII-5000`, `NSGAII-4750`, and `LBDF-250` on the DTLZ and WFG suites with 3 objectives. The Mann–Whitney U test with Bonferroni correction was applied to verify any statistically significant difference between the distributions of `NSGAII-5000` and `LBDF-250` (**** p-value ≤ 0.0001, *** p-value ≤ 0.001, ** p-value ≤ 0.01, * p-value ≤ 0.05, ns p-value ≤ 1).

In `NSGAII-5000` we used 50 individuals and 100 generations to balance the exploration and exploitation capabilities of NSGA-II. In `NSGAII-4750` we leveraged 95% of the budget to find a good quality first Pareto set, while keeping some budget to furtherly improve its quality with `LBDF-250`. Since we considered three-objective optimization problems, we used a different approach to determine pairs of adjacent solutions in the objective space. Namely, for each non-dominated solution, the nearest non-dominated solution (identified by the smallest Euclidean distance) in the objective space is selected, and an LBDF is applied between these points in the search space. In the case an LBDF was already applied between these two solutions, the second nearest non-dominated solution is selected. As before, we performed 30 runs to collect statistically significant tests.

Figure 6 shows the box-plot representation of the distributions of the hyper-volume values obtained in the three tests. According to these results, the approach based on LBDFs is better than `NSGA-II`, with a statistical significance difference on 8 out of 16 benchmark functions. It is worth noting that our approach improves the first Pareto set on many benchmark functions of the WGF suite, which includes more difficult problems than DTLZ [34].

We executed the same tests using NSGA-III with reference directions initialized according to the Riesz s-Energy method [39]. The obtained results showed that also in this case the approach based on LBDFs is significantly better than NSGA-III on 8 out of 16 benchmark functions (data not shown). In particular, similarly to the tests performed with NSGA-II, the approach exploiting LBDFs improves the first Pareto set of many benchmark functions of the WFG suite.

IV. Conclusions

In this work, we presented an approach, based on local-bounded dilations of the search space, designed to unveil non-dominated candidate solutions that were missing in the Pareto front. Namely, starting from an existing Pareto front we apply an LBDF, placed between each pair of candidate solutions that are adjacent in the objective space. We tested the effectiveness of our approach on Pareto fronts generated by using random sampling. Moreover, we also applied our approach to the Pareto fronts optimized by NSGA-II and NSGA-III. Our results show that our approach can improve the existing Pareto fronts using both strategies.

There are several directions in which this work could be further elaborated, the most important one is the generalization to problems with more than three objectives. When two-objectives problems are considered, the current algorithm leverages the strict ordering induced by the domination relationship to determine the adjacent non-dominated solutions of the first Pareto set. Conversely, with three-objectives problems, the concept of adjacency is defined as the closest non-dominated solution in terms of Euclidean distance. Although the results obtained with the three objectives are promising, we argue that the introduction of alternative definitions of adjacency based on geometrical properties might improve the performance of our approach when three (or more) objectives are considered. It is worth noting that more sophisticated definitions of adjacency can increase the computational costs or the number of LBDFs introduced, making such an approach unfeasible. Moreover, we speculate that using some adaptive strategy for the selection of the parameters γ and r of the LBDF might improve the outcome of our approach.

Finally, we also plan to consider additional metrics, for example the Pure Diversity Indicator [40], to evaluate other geometrical characteristics of the first Pareto front; e.g., how much the solutions are spread. We will investigate these improvements in our future work.

REFERENCES

[1] A. Kumar, G. Wu, M. Z. Ali, R. Mallipeddi, P. N. Suganthan, and S. Das, "A test-suite of non-convex constrained optimization problems from the real-world and some baseline results," *Swarm and Evolutionary Computation*, vol. 56, p. 100693, 2020.

[2] A. Tangherloni, S. Spolaor, P. Cazzaniga, D. Besozzi, L. Rundo, G. Mauri, and M. S. Nobile, "Biochemical parameter estimation *vs*. benchmark functions: A comparative study of optimization performance and representation design," *Applied Soft Computing*, vol. 81, p. 105494, 2019.

[3] K. Deb, *Optimization for engineering design: Algorithms and examples*. PHI Learning Pvt. Ltd., 2012.

[4] M. Hajihassani, D. J. Armaghani, and R. Kalatehjari, "Applications of particle swarm optimization in geotechnical engineering: a comprehensive review," *Geotechnical and Geological Engineering*, vol. 36, no. 2, pp. 705–722, 2018.

[5] S. Shishegar, S. Duchesne, and G. Pelletier, "Optimization methods applied to stormwater management problems: a review," *Urban Water Journal*, vol. 15, no. 3, pp. 276–286, 2018.

[6] M. Elbes, S. Alzubi, T. Kanan, A. Al-Fuqaha, and B. Hawashin, "A survey on particle swarm optimization with emphasis on engineering and network applications," *Evolutionary Intelligence*, vol. 12, no. 2, pp. 113–129, 2019.

[7] L. Rundo, A. Tangherloni, C. Militello, M. C. Gilardi, and G. Mauri, "Multimodal medical image registration using particle swarm optimization: A review," in *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE, 2016, pp. 1–8.

[8] F. Reali, C. Priami, and L. Marchetti, "Optimization algorithms for computational systems biology," *Frontiers in Applied Mathematics and Statistics*, vol. 3, p. 6, 2017.

[9] M. S. Nobile, A. Tangherloni, L. Rundo, S. Spolaor, D. Besozzi, G. Mauri, and P. Cazzaniga, "Computational intelligence for parameter estimation of biochemical systems," in *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE, 2018, pp. 1–8.

[10] A. Tangherloni, S. Spolaor, L. Rundo, M. S. Nobile, P. Cazzaniga, G. Mauri, P. Liò, I. Merelli, and D. Besozzi, "GenHap: a novel computational method based on genetic algorithms for haplotype assembly," *BMC Bioinformatics*, vol. 20, no. 4, p. 172, 2019.

[11] A. Tangherloni, S. G. Riva, S. Spolaor, D. Besozzi, M. S. Nobile, and P. Cazzaniga, "The impact of representation on the optimization of marker panels for single-cell RNA data," in *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE, 2021, pp. 1423–1430.

[12] A. Tangherloni, S. G Riva, B. Myers, and P. Cazzaniga, "Multi-objective optimization for marker panel identification in single-cell data," in *Proceedings of the IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology*. IEEE, 2022, pp. 1–8.

[13] K. Deb, "Multi-objective optimization," in *Search methodologies: Introductory tutorials in optimization and decision support techniques*. Springer, 2014, pp. 403–449.

[14] A. Zhou, B.-Y. Qu, H. Li, S.-Z. Zhao, P. N. Suganthan, and Q. Zhang, "Multiobjective evolutionary algorithms: A survey of the state of the art," *Swarm and Evolutionary Computation*, vol. 1, no. 1, pp. 32–49, 2011.

[15] E. Zitzler and L. Thiele, "Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach," *IEEE Transactions on Evolutionary Computation*, vol. 3, no. 4, pp. 257–271, 1999.

[16] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.

[17] E. Zitzler, M. Laumanns, and L. Thiele, "SPEA2: Improving the strength Pareto evolutionary algorithm," *TIK-report*, vol. 103, 2001.

[18] K. Deb and H. Jain, "An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints," *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 577–601, 2013.

[19] H. Jain and K. Deb, "An evolutionary many-objective optimization algorithm using reference-point based nondominated sorting approach, part ii: Handling constraints and extending to an adaptive approach," *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 602–622, 2013.

[20] J. Hu, Y. Wang, E. Zhou, M. C. Fu, and S. I. Marcus, "A survey of some model-based methods for global optimization," *Optimization, Control, and Applications of Stochastic Systems*, pp. 157–179, 2012.

[21] R. B. Gramacy, *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC, 2020.

[22] T. Lew, A. Spencer, F. Scarpa, K. Worden, A. Rutherford, and F. Hemez, "Identification of response surface models using genetic programming," *Mechanical Systems and Signal Processing*, vol. 20, no. 8, pp. 1819–1831, 2006.

[23] L. Manzoni, D. M. Papetti, P. Cazzaniga, S. Spolaor, G. Mauri, D. Besozzi, and M. S. Nobile, "Surfing on fitness landscapes: a boost on optimization by Fourier surrogate modeling," *Entropy*, vol. 22, no. 3, p. 285, 2020.

[24] F. Zhao, G. Yang, Y. Zhang, W. Lei, W. Ma, and C. Zhang, "A novel hybrid combination optimization algorithm based on search area segmentation and fast Fourier transform," *Engineering Optimization*, vol. 51, no. 5, pp. 846–861, 2019.

[25] J. G. Falcón-Cardona, E. Covantes Osuna, C. A. Coello Coello, and H. Ishibuchi, "On the utilization of pair-potential energy functions in multi-objective optimization," *Swarm and Evolutionary Computation.*, vol. 79, p. 101308, 2023.

[26] Y. Wang, Z. Cai, and Y. Zhou, "Accelerating adaptive trade-off model using shrinking space technique for constrained evolutionary optimization," *International Journal for Numerical Methods in Engineering*, vol. 77, no. 11, pp. 1501–1534, 2009.

[27] H. Wang, Z. Wu, Y. Liu, J. Wang, D. Jiang, and L. Chen, "Space transformation search: a new evolutionary technique," in *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 2009, pp. 537–544.

[28] F. Chunming, X. Yadong, C. Jiang, H. Xu, and Z. Huang, "Improved differential evolution with shrinking space technique for constrained optimization," *Chinese Journal of Mechanical Engineering*, vol. 30, no. 3, pp. 553–565, 2017.

[29] A. H. Aguirre, S. B. Rionda, C. A. Coello Coello, G. L. Lizárraga, and E. M. Montes, "Handling constraints using multiobjective optimization concepts," *International Journal for Numerical Methods in Engineering*, vol. 59, no. 15, pp. 1989–2017, 2004.

[30] M. S. Nobile, P. Cazzaniga, and D. A. Ashlock, "Dilation functions in global optimization," in *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE, 2019, pp. 2300–2307.

[31] D. M. Papetti, D. A. Ashlock, P. Cazzaniga, D. Besozzi, and M. S. Nobile, "If you can't beat it, squash it: Simplify global optimization by evolving dilation functions," in *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE, 2021, pp. 1414–1422.

[32] D. M. Papetti, V. Coelho, D. A. Ashlock, P. Cazzaniga, S. Spolaor, D. Besozzi, and M. S. Nobile, "Local bubble dilation functions: Hypersphere-bounded landscape deformations simplify global optimization," in *Proceedings of the IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology*, 2022, pp. 1–8.

[33] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, "Scalable test problems for evolutionary multiobjective optimization," in *Evolutionary Multiobjective Optimization*. Springer, 2005, pp. 105–145.

[34] S. Huband, P. Hingston, L. Barone, and L. While, "A review of multiobjective test problems and a scalable test problem toolkit," *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 5, pp. 477–506, 2006.

[35] K. Shang, H. Ishibuchi, L. He, and L. M. Pang, "A survey on the hypervolume indicator in evolutionary multiobjective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 25, no. 1, pp. 1–20, 2021.

[36] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *Annals of Mathematical Statistics*, pp. 50–60, 1947.

[37] F. Wilcoxon, "Individual comparisons by ranking methods," in *Breakthroughs in Statistics*. Springer, 1992, pp. 196–202.

[38] O. J. Dunn, "Multiple comparisons among means," *Journal of the American statistical association*, vol. 56, no. 293, pp. 52–64, 1961.

[39] J. G. Falcón-Cardona, H. Ishibuchi, and C. A. Coello Coello, "Riesz s-energy-based reference sets for multi-objective optimization," in *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE, 2020, pp. 1–8.

[40] H. Wang, Y. Jin, and X. Yao, "Diversity assessment in many-objective optimization," *IEEE Transactions on Cybernetics*, vol. 47, no. 6, pp. 1510–1522, 2016.