

A Novel Population-based Multi-Objective CMA-ES and the Impact of Different Constraint Handling Techniques¹

Sílvia Rodrigues ^a Pavol Bauer ^a Peter A.N. Bosman ^b

^a *Delft University of Technology (TU Delft), Delft, The Netherlands*

^b *Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands*

1 Introduction

Many real-world problems have more than one conflicting objective that need to be optimized simultaneously. Moreover, many problems require taking a black-box optimization (BBO) perspective, i.e. assume that (virtually) nothing is known (e.g. complex simulation-based real-world models). Studying and understanding algorithms to tackle optimization problems under such conditions is therefore important.

The Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) is a well-known, state-of-the-art optimization algorithm for single-objective real-valued (BBO) problems. Although several extensions of CMA-ES to multi-objective (MO) optimization exist, none incorporates a key component of the most robust CMA-ES variant: associating a population with each Gaussian that drives optimization.

Many real-world problems also have constraints, making the performance of BBO algorithms under different constraint-handling techniques important. All MO-CMA-ES variants previously introduced use a penalty term to handle box constraints. Although this resulted in fast convergence speeds for certain benchmark problems, it also has drawbacks. For example, it only performs well with box constraints since these, differently from general problem constraints, allow an easy mapping to the feasible space. Furthermore, infeasible solutions may end up in the elitist archive.

The main objectives of this paper stem from the fact that all existing MO-CMA-ES variants use populations of size one and that only the penalty approach was used to handle constraints. Our first goal is to study the benefits of having a population-based MO-CMA-ES. To do so, we study a combination between a multi-objective optimization framework that was recently introduced [1] with the most general SO version of CMA-ES [2]. Our second goal is to assess the performance and robustness of the previously introduced MO-CMA-ES variants, the novel population-based MO-CMA-ES and the iMAMaLGaM algorithm [1] under different and more general constraint handling techniques.

2 Population-based Multi-Objective CMA-ES

A general framework for extending population-based algorithms from single- to multi-objective optimization was introduced in [1]. We made a few improvements to this framework and used it to construct a population-based multi-objective version of CMA-ES. For the minor improvements, we refer the reader to the full paper. Here we only give a high-level flavor of the workings of the framework. In addition to common domination-rank-based selection, variation is ensured to be based on clustering. That is, the selected solutions are explicitly clustered in the objective space and each cluster undergoes variation separately. An important part of state-of-the-art variation operators are adaptive mechanisms that span multiple generations. The performance of these mechanisms strongly depends on a correlation between the solution sets in subsequent generations. Therefore, some form of registration is required to

¹The full paper has been accepted for presentation at The Genetic and Evolutionary Computation Conference (GECCO 2014).

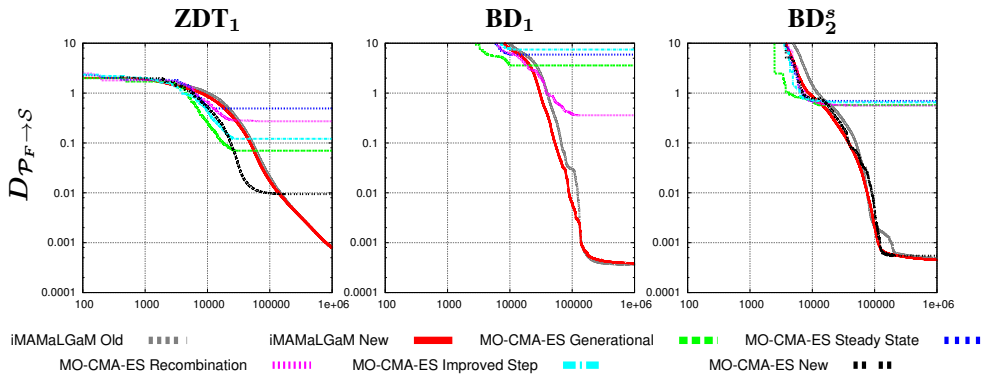


Figure 1: Performance of all algorithms on selected problems, averaged over 100 runs. Horizontal: number of evaluations. Vertical: distance-to-optimum indicator (lower is better, 0.001 is value to reach).

determine the best correspondence between clusters in subsequent generations, for which a greedy algorithm is used to construct the best pairing of clusters in subsequent generations. Every cluster is then allowed to generate an equal number of solutions through variation whereby for each objective one cluster is identified that focuses solely on making improvements in that objective so as to specifically add pressure on extending the Pareto front along a single axis. Finally, an elitist archive is maintained with all currently non-dominated solutions. If the objectives are real valued, infinitely many non-dominated solutions are possible. To prevent the archive from growing to extreme sizes, the objective space is adaptively discretized into hypercubes so as to accommodate a predefined desired number of solutions.

3 Results

The results are averaged over 100 independent runs and shown for the most interesting cases in Figure 1 (for more graphs, see the full version of this paper). The differences between the old and new versions of iMAMaLGaM have a positive, but small effect in terms of convergence.

No MO-CMA-ES from literature could solve unconstrained problem BD_2^s . Only one of the Pareto extremes was found because BD_2^s has one objective that is far simpler than the other, resulting in the population being pulled quickly toward one end of the Pareto front. Due to the lack of pressure toward improving individual objectives as in the MO framework, it is hard for existing MO-CMA-ES implementations to find the other Pareto extreme, resulting in very low convergence speed and early stopping the optimization runs. Population-based MO-CMA-ES and the iMAMaLGaM variants did solve BD_2^s .

The other problems do have constraints, making them harder, especially for CMA-ES. Although the population-based variant of CMA-ES always did better than existing CMA-ES variants, no variant of CMA-ES could solve the constrained problems satisfactorily using any tested constraint-handling technique. This is due to the careful way CMA-ES has been designed. When more non-smooth search-space adaptations occur, for instance as a result of constraint handling, the powerful search-space exploitation capacity of CMA-ES breaks down. In this regard, the iMAMaLGaM basis appears to be more robust as its multi-objective counterpart could solve all problems with all constraint handling techniques.

4 Conclusions

We introduced a novel population-based MO-CMA-ES. Experimental results demonstrate that the proposed approach is, in general, more robust when compared to the multi-objective extensions of CMA-ES that were previously introduced in literature. Furthermore, the algorithms based on CMA-ES demonstrated to be very sensitive to the way constraints were handled. By comparison, iMAMaLGaM showed to be more robust since it was able to solve all problems with all constraint handling techniques.

References

- [1] Peter A. N. Bosman and T. Alderliesten. Incremental Gaussian model-building in multi-objective edas with an application to deformable image registration. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO-2012*, pages 241–248, New York, 2012. ACM Press.
- [2] N. Hansen and S. Kern. Evaluating the CMA evolution strategy on multimodal test functions. In X. Yao et al., editors, *Parallel Problem Solving from Nature (PPSN VIII)*, pages 282–291. Springer, 2004.