SS 2021-2022, SMT-ROB, DISAL-SP171 Project Summary

Improving Mixed-Discrete Particle Swarm Optimization for Categorical Problems

Léonard Pasi

Professor : Alcherio Martinoli
Assistant(s) : Cyrill Baumann

Particle Swarm Optimization (PSO) is a population-based metaheuristic algorithm. As such, a pool of candidate solutions (i.e., a “swarm of particles”) moves across the optimization space in search of the optimum, and no assumption on the problem is needed. Since its introduction in 1995 for continuous optimization, it has enjoyed a lot of success, and numerous variants have been proposed. Some of them aimed to make PSO suitable for mixed-discrete problems. For instance, Chowdhury et al. proposed a Mixed Discrete PSO (MDPSO) in 2013. One limitation of MDPSO is that in principle, it cannot deal with categorical variables, although in practice it still can, holding decent performance. Note that in general discrete variables are either ordinal or categorical, and problems that include categorical variables are called Mixed-Variable Optimization Problems (MVOPs). The goal of this project was to adapt PSO for MVOPs, implement it in C++ and compare it with the existing MDPSO implementation.

A literature search revealed that a PSO variant for MVOPs, called PSO_{mv}, was in fact proposed in 2021 by Wang et al, and showed very competitive performance. PSO_{mv} treats continuous and discrete variables separately, while no distinction is made between ordinal and categorical variables. The update techniques for the continuous and discrete parts of the particles are both novel. Finally, thanks to an adaptive parameter tuning strategy, PSOmv has (almost) no hyper-parameter. Regrettably, PSO_{mv} was only tested as a whole, thus the individual merits of its three main components remained unclear. This posed a significant hurdle to investigate potential improvements of the algorithm. Additionally, the overall scientific rigor of the validation could be ameliorated.

Scientifically sound performance assessment is critical to fairly compare optimization algorithms. Unfortunately, it is also a tedious and repetitive task. In this context, the use of standardized tools, such as COCO, comes as a logical choice. COCO is an open-source platform for Comparing Continuous Optimizers in a black-box setting (in fact, Mixed-Discrete optimizers can also be compared). It allows to both conduct experiments and process automatically the gathered data. Its main performance metric, based on runtimes (measured in number of objective function evaluations) to reach one or several fitness target values, is the only available measure with a generic, meaningful, and quantitative interpretation.

After implementing PSO_{mv} in C++, and interfacing it with the COCO platform, the algorithm was tested extensively on both mixed-integer and continuous problems. It turned out that the adaptive parameter tuning strategy and the novel technique for continuous variable did not bring any significant benefit for the problems considered. Overall, MDPSO and PSO_{mv} have similar performance on mixed-integer problems.

However, PSO_{mv} still has one important advantage: it is naturally suitable for MVOPs. Since MDPSO can also be used for MVOPs, this advantage had to be confirmed in practice. As COCO doesn’t support benchmarking on MVOPs, a small set of MVOPs coded by DISAL was used. The generated plots, such as the one shown below, suggest that PSO_{mv} is in fact outperforming MDPSO on MVOPs. However, more extensive evaluation is needed for a conclusive conclusion.

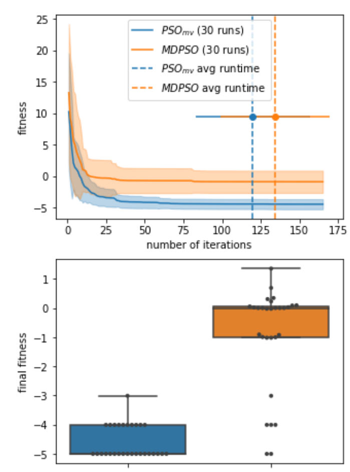

Figure 1: Comparison between MDPSO and PSOmv on one MVOP.