Diversity Oriented Local Search for Multi-criteria Optimization

Oded Maler Abhinav Srivastav

April 8, 2015

Abstract

We develop a generic tool for approximating the Pareto front of multi-criteria optimization problems using stochastic local search algorithms. Our algorithmic scheme handles problems that the multi-criteria context introduces into the local search framework such as the non-uniqueness of the best neighbor and the potentially large size of the Pareto front. We demonstrate the performance of our algorithm under different configurations and parameters on multi-criteria variants of the quadratic assignment and 0-1 knapsack problems. The adaptation of our scheme to new problems involves a minimal investment.

1 Introduction

The present work is a contribution toward advancing the state-of-the-art in computational technique for approximating Pareto fronts of multi-criteria optimization (MCO) problems. More specifically, we develop a generic stochastic local search algorithm adapted to MCO. Local search is a powerful heuristic technique for solving difficult combinatorial optimization problems. The basic step in local search it to compute the neighborhood of a given solution, the solutions which can be obtained from it by a single modification and selecting with high probability the neighbor with the optimal cost. In MCO, it is often the case that there is a set of non-dominated and mutually comparable neighbors. Any local search algorithm must choose one of these neighbors for exploration in the next iteration. Furthermore, there is an additional question of discarding or keeping the unselected candidate solutions. Our strategy maintains such solutions and can select them in subsequent steps. Another major issue concerning the MCO problem is the number of non-dominated solutions. In many real-life situations, the number of non-dominated solutions can be prohibitively large. Keeping such a large list of solutions as well as large list of candidate solutions may render the algorithm computationally inefficient. This is primarily due the fact that the local search algorithm has to compare the cost of newly found solutions with those already present; the process often referred to as Pareto filtering. Furthermore, providing a very large set of solutions is not always productive as a decision making aid.

2 Contribution

Our algorithm employs a new parameterized diversity based reduction technique to maintain a small set of solutions and candidate solutions. This new diversity measure D(c,P) estimates the contribution of a point c to the diversity of set of solutions in P. Intuitively, the more c is isolated from other points in P, the larger is D(c,p). The salient features of our algorithm are (1) A selection policy based on diversity-oriented measure; (2) Keeping track of minimal points that have not been selected and including them as candidate for future iterations; (3) Update policies that keep the sets of optimal and candidate solutions small.

We developed a generic tool for extending local search to MCO problems. The tool is written in a modular way where the specifics of the particular problem to be solved are separated from the ingredients and parameters of the local search algorithm. Hence, to apply the tool to a new problem it suffices to write some small code that specifies how costs are derived from a representation of a solution and how local move operators are defined.

3 Experiments

For analysis, we applied our tool extensively to two classes of problems: the bi-objective 0/1 knapsack problem (bKNAP) and the bi-objective quadratic assignment problems (bQAP). For the bKNAP problem, we consider several groups of instances; uncorrelated instance, weakly correlated instance, strongly correlated instance and subset sum instance with large number of items. For bQAP problem we consider uncorrelated and correlated instances with large number of facilities. Pareto sets are compared using the hyper-volume unary indicator. For each instance, the reference set is obtained by comparing all the runs for that instance. Therefore, the outperformance of an algorithm A on some other algorithm B is based on the simple dominance of volume of A on B. First, we compare our selection procedure with the strategy that randomly selects solutions from local Pareto set (non-dominated solutions in the neighborhood). These runs are performed with unbounded set of solutions along with unbounded set of candidate solutions. We show that for the most of the instances of bKNAP problem, our method outperforms random selection. In case of bQAP problems, our method performs similar to random selection procedure for small instances with number of facilities equal to 25, but as the size of the problem increases, our selection outperforms random selection. The experiments suggest that random strategy is more biased towards the convex inner portions of the objective space, where it outperforms our selection strategy. On other hand, our selection strategy results in better extreme solutions. Additionally, the number of approximate Pareto solutions found by our strategy is more than the number of solutions found by a random strategy. And this difference increases with the size of the problem. Next, we compare our strategy with volume based selection where the preference is given to those candidate solutions which have larger dominated volume.

As the number of non-dominated solutions increases with the size of the problem, we turn to bounded set for solutions and candidate solutions. We compare our strategy with random strategy and clustering based techniques. In random strategy, a solution is randomly discarded for the set of solutions such that the number of solutions remains bounded. In a clustering strategy, the solutions which are farthest from its corresponding cluster centroid are discarded. We show that our diversity based comparison outperforms both strategies using volume indicator and Hausdorff distance. We also show experimentally the effect of different thresholds on the size of the set of non-dominated solutions candidate solutions.

4 Conclusion

To conclude, we developed a powerful framework for solving multi-objective combinatorial optimization problems. We tested it extensively for QAP and knapsack problems and it can be easily applied to other types of problems. We also introduced a new diversity based measure which can be used in combination with other variants of local search algorithms as selection and archiving strategy. Additionally, it can also be employed as fitness function and selection strategy for mating individuals in evolutionary algorithms.