Toward the Design of Efficient Pivoting Rules for Local Search

Sara Tari
LERIA, Universite d'Angers, France
sara.tari@univ-angers.fr

Matthieu Basseur
LERIA, Universite d'Angers, France
matthieu.basseur@univ-angers.fr

Adrien Goëffon
LERIA, Universite d'Angers, France
adrien.goeffon@univ-angers.fr

1. INTRODUCTION

In the context of local search and intensification mechanisms, we study here move selection principles that allow the attainment of high-quality local optima while restricting the scope to hill-climbings. Previous study [1] showed that among classical pivoting rules, first improvement outperforms best improvement on landscapes recognized as difficult to climb (rugged and/or large). A further work [2] indicated that worst improvement is more efficient to attain good local optima than classical pivoting rules on difficult landscapes. An interesting observation is that the efficiency of worst improvement resides in its capacity to keep many improvement possibilities during the hill-climbing process.

These results drove us to design a pivoting rule based on this principle, attempting to reduce the probability to be trapped prematurely in low quality local optima. To achieve this, we propose the *maximum expansion* pivoting rule, which consists of selecting the improving neighbor which have the maximal number of improving neighbors.

2. FITNESS LANDSCAPES

A *Fitness Landscape* is a triplet $(\mathcal{X}, \mathcal{N}, f)$ where $\mathcal{X}$ is a set of solutions called *search space*, $\mathcal{N}$ a relation which associates to any solution a set of neighbors, and f a fitness function which assigns a scalar value to each solution.

Characteristics of fitness landscapes can affect significantly the behavior of search algorithms, therefore analyzing the links between fitness landscapes properties and algorithms efficiency can improve the understanding of evolutionary components. The difficulty to explore a landscape resides in particular in its size, amount of neutrality, and its ruggedness which is directly related to the number of local optima of the landscape [5]. Artificial combinatorial landscapes with tunable size and ruggedness can be generated using the NK model [4]. Size and ruggedness of NK landscapes are respectively tunabled by means of parameters N and K.

Climbing a combinatorial fitness landscape $(\mathcal{X}, \mathcal{N}, f)$ from an initial solution $x \in \mathcal{X}$ consists of iteratively choosing improving neighbors thanks to a pivoting rule, until a local optimum is reached (i.e. a solution x s.t. $\forall x' \in \mathcal{N}(x), f(x') < f(x)$). Consequently, the capacity of the search to reach good local optima depends on the designed pivoting rule as well as the landscape's characteristics [3].

Although there exists many ways to conceive pivoting rules, in most cases classical strategies *first* and *best* improvement are effectively used as basis of hill-climbing components within metaheuristics. Recall that first improvement selects the first improving neighbor encountered while best improvement evaluates the whole neighborhood at each step of the search and selects the improving one which has the highest fitness value. In a same way, we also consider the *worst* improvement rule [2], which generates the whole neighborhood at each step and selects the improving neighbor having the lowest fitness value.

3. MAXIMUM EXPANSION

Experiments conducted in [2] highlighted that worst improvement often allows the attainment of good quality local optima although the search requires more steps. This indicates that such a pivoting rule tends to delay the local optimum attainment. Thus, we intuitively deduced that a specificity of worst improvement is to maintain a high number of improving possibilities during the search. It makes sense to conjecture that the efficiency of worst improvement on rugged landscapes resides essentially in its capacity to keep a high number of improvement possibilities during the search process. This observation led us to propose the *maximum expansion* pivoting rule (ME) which aims at keeping the most improvement possibilities during the search.

ME consists of selecting the improving neighbor which possess the higher number of improving neighbors (*expansion score*). When several improving neighbors have the same non-null expansion score, one of them is randomly selected. Once all improving neighbors have their expansion scores equal to zero, ME selects the best improving neighbor since this necessarily corresponds to the final step of the climbing process.

We first focus on the quality of local optima reached by classical pivoting rules compared to those obtained by ME. However, let us notice that ME evaluates several neighborhoods at each step, which lead to significantly increase the number of evaluation compared to commonly used moving strategies.

Experiments are realized on NK landscapes in order to assess ME ability to reach good local optima with respect to the characteristics of landscapes. We use a benchmark

GECCO '16, July 20–24, 2016, Denver, Colorado, USA.

ACM ISBN 978-1-4503-4206-3/16/07..

DOI: http://dx.doi.org/10.1145/2908961.2908993

set of 280 NK landscapes of various log-sizes and variable interdependency levels.

We conduct 100 hill-climbing executions for each couple (pivoting rule, landscape) starting from identical sets of 100 randomly generated initial solutions. For $K = 1$, ME is often statistically dominated by other pivoting rules. This result confirms that the best improvement pivoting rule is clearly appropriated for climbing very smooth landscapes. Except for this specific case, ME always statistically dominates the three other considered pivoting rules. This notably includes rugged landscapes, which are known to be more challenging to tackle using local search algorithms.

Although ME is particularly efficient to attain good local optima, its computational cost cannot compete with best, and even worst improvement which is already time-consuming. Let us precise that although worst improvement requires more steps compared to ME, the latter evaluates more solutions at each step of the search. First improvement avoids the generation of the whole neighborhood and drastically reduces the number of evaluations required to attain a local optimum. In a previous work [2], the worst improvement strategy was proposed in a approximated version which stops the neighborhood evaluation when a predefined number of improving neighbors has been encountered.

4. ME APPROXIMATION

In order to reduce the computional cost of ME, we propose investigate three approximation variants: ME_k, ME^l and ME_k^l. The first variant ME_k, which is directly derived from the worst improvement approximation, works as follow. First, we evaluate the neighborhood until at most k improving neighbors have been encountered. Among them, the neighbor with the best expansion score is selected. Thus ME_k requires to evaluate at most k complete neighborhoods at each step of the search. Indeed the current neighborhood only needs to be (partially) evaluated at the first step of the search since neighborhoods of selected neighbors are already known in ME searches. Consequently, parameter k directly affects the number of evaluations per step.

Since the number of evaluations required for ME mainly comes from the calculation of the expansion score, we propose in ME^l to only approximate the determination of an improving neighbor with the maximal expansion score. Neighborhoods of all improving neighbors are evaluated simultaneously; the evaluation stops once a neighbor reaches an expansion score of l. When the maximum expansion score is inferior to l for each neighbor, ME^l simply selects at random an improving neighbor with the best expansion score.

It is obvious that approximation techniques provided by ME_k and ME^l are not able to reduce sufficiently the computational cost induced by ME, since each step requires the evaluation of at least one complete neighborhood. Then we propose ME_k^l which is a combination of ME_k and ME^l. At each step, k improving neighbors are considered and ME_k^l selects the one having l improving neighbors while requiring the minimal number of neighbors evaluation.

Results show that for each type of landscapes tested, ME_k^* outperforms first, best and worst improvement when $k > 4$ on rugged landscapes. The approximation of the current solution neighborhood performed by ME_k^* provides local optima which remain better than those attained by classical pivoting rules, at least when ME statistically dominates usual climbers ($K > 1$).

ME_*^l dominates most of the classical pivoting rules for $k \geq 8$. Considering a fixed value of l, the precision of the expansion score approximation decreases while N increases; ME_*^l becomes consequently less efficient for larger landscapes with a given value of k. ME_*^l requires relatively less evaluations but reaches lower local optima than ME_k^* when $k = l$. It makes sense that ME_*^l approximation is less efficient since the expansion score is approximated while ME_k^* approximates the direct neighborhood.

The computational cost of the two considered approximation methods is drastically reduced in comparison to ME and is comparable to worst improvement, especially for large and rugged landscapes. However, worst improvement is time-consuming compared to the other two classical pivoting rules. Approximating ME with only setting l or k is not sufficient to reach an evaluation cost comparable to best improvement.

ME_k^l reaches good local optima and in a significantly reduced computational time in comparison to ME and other ME approximations. When $l > k$, better local optima are reached in an almost similar number of evaluations. For instance, ME_4^8 reaches higher local optima than ME_8^4. This approximation is more efficient when $N = 256$ than when $N = 1024$ which makes sense as l is fixed regardless to N. It produces a less accurate expansion score approximation for larger neighborhoods.

While ME allows the search to reach better local optima than classical pivoting rules on most NK landscapes, experiments of approximated variants indicate that good tradeoffs between efficiency (in terms of solution quality) and computational effort can be provided. Nevertheless, let us recall that the main aim of this work is to identify neighborhood search rules able to drive the search toward more promising areas. Such information may help to understand and anticipate local searches behavior, as well as to provide novel guidelines for elaborating efficient metaheuristics.

Acknowledgements

The work is partially supported by the PGMO project from the *Fondation Mathématique Jacques Hadamard*.

5. REFERENCES

[1] Matthieu Basseur and Adrien Goëffon. Hill-climbing strategies on various landscapes: an empirical comparison. In *Genetic and Evolutionary Computation Conference (GECCO)*, pages 479–486. ACM, 2013.

[2] Matthieu Basseur and Adrien Goëffon. On the efficiency of worst improvement for climbing NK landscapes. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, GECCO '14, pages 413–420, New York, NY, USA, 2014. ACM.

[3] Matthieu Basseur and Adrien Goëffon. Climbing combinatorial fitness landscapes. *Applied Soft Computing*, 30:688–704, 2015.

[4] Stuart A. Kauffman and Edward D. Weinberger. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology*, 141(2):211 – 245, 1989.

[5] Katherine M. Malan and Andries P. Engelbrecht. A survey of techniques for characterising fitness landscapes and some possible ways forward. *Information Sciences*, 241:148–163, 2013.