Exploring non-neutral Landscapes with neutrality-based Local Search

Matthieu Basseur, Adrien Goëffon, and Hugo Traverson

LERIA, Université d'Angers (France)
matthieu.basseur@univ-angers.fr
adrien.goeffon@univ-angers.fr
hugo.traverson@univ-angers.fr

Abstract. In this paper, we present a generic local search algorithm which artificially adds neutrality in search landscapes by discretizing the evaluation function. Some experiments on NK landscapes show that an adaptive discretization is useful to reach high local optima and to launch diversifications automatically. We believe that a hill-climbing using such an adaptive evaluation function could be more appropriated than a classical iterated local search mechanism.

1 Context

In combinatorial optimization, fitness landscapes study abstracts problem specificities and aims at evaluating the pertinence of generic metaheuristics. More formally, a fitness landscape is a triplet $(\mathcal{X}, \mathcal{N}, f)$, where $\mathcal{X}$ is a discrete set of solutions, $\mathcal{N} : \mathcal{X} \to 2^{\mathcal{X}}$ a neighborhood function, and $f : \mathcal{X} \to [0, 1)$ a fitness (evaluation) function. In [1], we compared the efficiency of several hill-climbing variants (denoted as *climbers*) to determine neighborhood-based moving strategies that are likely to reach *high* solutions (with high fitness values). In particular, we focused on the ways to handle neutrality [4]. To achieve this, we introduced rounded landscapes, by setting a discretization level of the fitness function, referring to its codomain size r. The rounded function f_r is then defined from an original fitness function f as follows:

$$f_r(x) = \frac{\lfloor r.f(x) \rfloor}{r}$$

Let us notice that f_r gives a partial order which is compliant with the order relation induced by f. It means that, $\forall x, y \in \mathcal{X}, f(x) < f(y) \Rightarrow f_r(x) \leqslant f_r(y)$. A second property of f_r functions ($f_r(x) \leqslant f(x) < f_r(x) + \frac{1}{r}$) makes possible to compare fitnesses reached on original and rounded corresponding landscapes. In [1], we observed on rounded landscapes that some r values allow *stochastic climbers* (first improvement which accepts indifferently neutral and improving moves) to reach higher solutions than while considering original fitness functions.

In this paper, we propose to extend the principle of rounding fitness function to help neutrality-based local searches to reach high local optima. We propose a generic local search algorithm based on an adaptive evaluation function $LS_{\overline{f}}$ which incorporates an artificial rate of neutrality chosen with respect to information collected during the search. It simulates deteriorating moves during intensification phases, as well as automatic perturbations when local improvements become rare.

Here, the efficiency of the proposed mechanism is evaluated on NK landscapes [2], where size and ruggedness are tunable by means of parameters N and K.

2 Determining appropriate neutrality rates of landscapes

Intuitively, adding neutrality to landscapes necessarily induces to decrease their ruggedness, which make local searches more efficient as long as they exploit neutral moves. Nevertheless, a too large neutrality level creates flat areas which can drastically increase the number of moves needed to reach high solutions. In an extreme case, on totally flat landscapes, stochastic hill-climbings behave like random walk processes. Here, we propose to control the neutrality rate by means of the fitness function f_r. Setting r adequately consists in reducing the ruggedness while keeping a moderate rate of neutrality.

We have extended our previous study on NK and NKr landscapes in order to determine the most appropriate rounding values with respect to landscapes under consideration (with $N \in \{128, 256, 512, 1024\}$ and $K \in \{1, 2, 4, 8\}$). To estimate the neutrality of these most appropriate rounded landscapes, we define a search neutrality indicator $\tilde{\nu}$ which aims to estimate the average rate of neutrality encountered during the search:

Definition 1 (Search neutrality). *Let $\mathcal{P}$ be a fitness landscape, N the neighborhood size and C the history of a climber execution (given by a sequence of evaluated and selected solutions). Let p the number of strictly improving moves in C. l_i and n_i ($i \in \{1, \ldots, p\}$) refer respectively to the number of evaluations and the number of neutral moves realized between the $(i-1)^{th}$ and i^{th} improving moves. The search neutrality depends on C and P, and is defined by:*

$$\tilde{\nu}(C, \mathcal{P}) = \frac{\sum_{i=1}^{p} \frac{n_i}{l_i} \log_N l_i}{\sum_{i=1}^{p} \log_N l_i}$$

Table 1 reports the ranges of rounding values r from which climbers executed on corresponding NKr landscapes are not statistically outperformed by climbers executed on derived landscapes with other values of r, as well as their associated search neutrality. Non-dominated ranges have been determined using a dichotomic sampling of r values. Statistical analysis were performed using a binomial test based on 100 runs per r value. We observe that optimal r values depend on landscape properties. The most noticeable information is that the $\tilde{\nu}$ values which make a neutrality-based climbing more efficient, are similar on all (N,K) parameterizations.

N, K	NK	NKr (NK using f_r)		
	f mean		r range	$\tilde{\nu}$ range
128_1	.7021	.7227	$[206, 285]$	$[20, 26]$
128_2	.7021	.7395	$[94, 148]$	$[23, 32]$
128_4	.7254	.7837	$[54, 81]$	$[23, 31]$
128_8	.7142	.7685	$[40, 64]$	$[19, 25]$
256_1	.7021	.7206	$[321, 350]$	$[26, 28]$
256_2	.7066	7410	$[212, 264]$	$[24, 28]$
256_4	.7235	.7843	$[117, 152]$	$[24, 28]$
256_8	.7166	.7755	$[84, 112]$	$[19, 22]$

N, K	NK	NKr (NK using f_r)		
	f mean		r range	$\tilde{\nu}$ range
512_1	.6897	.7069	$[731, 824]$	$[24, 27]$
512_2	.7135	.7484	$[454, 523]$	$[25, 28]$
512_4	.7200	.7790	$[234, 304]$	$[24, 28]$
512_8	.7206	.7810	$[173, 206]$	$[19, 22]$
1024_1	.6969	.7157	$[1371, 1562]$	$[25, 28]$
1024_2	.7146	.7507	$[959, 1178]$	$[23, 27]$
1024_4	.7246	.7844	$[449, 570]$	$[24, 28]$
1024_8	.7216	.7836	$[331, 403]$	$[19, 21]$

Table 1. Climber comparison on NK and non-dominated NKr landscapes. We also indicate ranges of optimal values of r and their associated $\tilde{\nu}$ (in %).

Input : a fitness landscape $(\mathcal{X}, \mathcal{N}, \mathfrak{f})$, a set of n rounding values $\mathcal{R} = \{r_1, \ldots, r_n\}$ (with $r_i > r_{i+1}$), parameters d, ν_{ref}, θ, ϵ.
Output: the best solution found x_{opt}
Randomly select $x \in \mathcal{X}$;
$F \leftarrow \mathfrak{f}(x)$;
$t \leftarrow 0$;
$F_{\text{opt}} \leftarrow F$;
$x_{\text{opt}} \leftarrow x$;
repeat
  Randomly select $x' \in \mathcal{N}(x)$;
  $t \leftarrow t + 1$;
  $\alpha \leftarrow 1 - \frac{1}{\min(d,t)}$;
  for all $r_i \in \mathcal{R}$ **do**
    $\nu_{\text{est}}[r_i] \leftarrow \alpha \times \nu_{\text{est}}[r_i]$;
    $\mu_{\text{est}}[r_i] \leftarrow \alpha \times \mu_{\text{est}}[r_i]$;
    if $f_{r_i}(x') = f_{r_i}(x)$ **then** $\nu_{\text{est}}[r_i] \leftarrow \nu_{\text{est}}[r_i] + 1 - \alpha$;
    else if $f_{r_i}(x') > f_{r_i}(x)$ **then** $\mu_{\text{est}}[r_i] \leftarrow \mu_{\text{est}}[r_i] + 1 - \alpha$;
  $\mathcal{R}' \leftarrow \{r_i \in \mathcal{R}, \nu_{\text{est}}[r_i] > 0\}$;
  if $t > d$ **then** $\mathcal{R}' \leftarrow \{r_i \in \mathcal{R}', \mu_{\text{est}}[r_i] \geqslant \epsilon\}$;
  if $\mathcal{R}' = \emptyset$ **then** $t \leftarrow 0$ {*Diversification*} ;
  else
    $R \leftarrow \operatorname{argmin}_{r_i \in \mathcal{R}'} \frac{\max(\nu_{\text{est}}[r_i], \nu_{\text{ref}})}{\min(\nu_{\text{est}}[r_i], \nu_{\text{ref}})}$;
    $\mathcal{R}' \leftarrow \{r_i \in \mathcal{R}',\ \nu_{\text{est}}[R] - \theta \leqslant \nu_{\text{est}}[r_i] \leqslant \nu_{\text{est}}[R]\}$;
    $R \leftarrow \operatorname{argmin}_{r_i \in \mathcal{R}'} \nu_{\text{est}}[r_i]$;
    if $f_R(x') \geqslant f_R(x)$ **then**
      $x \leftarrow x'$;
      if $f(x) > f(x_{\text{opt}})$ **then**
        $x_{\text{opt}} \leftarrow x$;
until *Stopping criterion*;

Algorithm 1: $LS_{\bar{f}}$

3 Local search with adaptive evaluation function ($LS_{\bar{f}}$)

Previous results showed that there exists an appropriate range of neutrality rate which leads to reach high solutions by hill-climbing. About 25% of neutrality seems appropriate for climbing efficiently NK landscapes. However, additional experiments emphasized that the local neutrality in NKr landscapes is negatively correlated with the height (fitness value) of solutions. Then, we propose to set r dynamically thanks to an adaptive mechanism which aims at preserving a reference neutrality rate ν_{ref}.

The local search algorithm we introduce here, $LS_{\bar{f}}$, is a stochastic climber which selects at each iteration a rounding value r_i among a set of candidate roundings $\{r_1, \dots, r_n\}$. To each rounding value r_i is associated an estimated neutrality $\nu_{\text{est}}[r_i]$, which is dynamically updated at each iteration as follows:

$$\nu_{\text{est}}[r_i] \leftarrow \begin{cases} \alpha \times \nu_{\text{est}}[r_i] + 1 - \alpha, \text{ if } f_{r_i}(x') = f_{r_i}(x) \\ \alpha \times \nu_{\text{est}}[r_i], \text{ otherwise} \end{cases}$$

Additionally, at each iteration, we select the rounding value which is the closest to ν_{ref} in terms of ratio (more precisely $\text{argmin}_{r_i} \frac{\max(\nu_{\text{est}}(r_i), \nu_{\text{ref}})}{\min(\nu_{\text{est}}(r_i), \nu_{\text{ref}})}$).

Maintaining a certain level of neutrality can prevent to reach local optima (in the sense of the original fitness function) and also requires to use a predefined number of iterations as a stopping criterion. Therefore, we propose to associate an improving move rate estimation μ_{est} to each rounding value r_i. $\mu_{\text{est}}[r_i]$ is estimated similarly to $\nu_{\text{est}}[r_i]$, by considering strictly improving moves. This allows the detection of search stagnation with respect to each rounding value. Then we refine the r_i selection mechanism by forbidding *stagnant* r_i values to be selected (r_i such that $\mu_{\text{est}}[r_i]$ is lower than a threshold ϵ). As a consequence, the search will be naturally driven to a local optimum.

To simulate a perturbation mechanism, we just need to reset the neutrality estimations. This can be done when every μ_{est} value, which estimate improving move rates, is smaller than a threshold ϵ. Such mechanism partially randomizes the search during several steps by considering flat landscapes. Algorithm 1 provides a detailed description of $LS_{\bar{f}}$.

To assess the relevance of the proposed climbing technique, $LS_{\bar{f}}$ has been compared with a classical Iterated Local Search (ILS) process, where diversification has been parameterized as follows:

- random restart ;
- random walk from the last local optimum found (5 variants: 5, 10, 15, 20, 30 moves) ;
- random walk from the best local optimum found (5 variants: 5, 10, 15, 20, 30 moves).

On each instance, $LS_{\bar{f}}$ is compared to the 11 ILS parameterizations (with 200 perturbations for each).

Table 2 compares average fitnesses reached by $LS_{\bar{f}}$ (with and without perturbations), with local searches (with — LS — and without perturbation —

ILS) using original evaluation functions, requiring an equivalent computational effort. These results emphasize the relevance of adapting the shape of a landscape according to its local properties. Moreover, $LS_{\bar{f}}$ allows the simulation of deteriorating moves during intensification and diversification phases without explicitly dealing with them. It is obvious that the behavior of such a mechanism can be linked with Simulated Annealing (SA) [3]. In future work, it should be interesting to provide a deep analysis of these two ways to simulate diversification during an intensification process.

Instance	LS	$LS_{\bar{f}}$ (0)	$LS_{\bar{f}}$ (200)	ILS	Instance	LS	$LS_{\bar{f}}$ (0)	$LS_{\bar{f}}$ (200)	ILS
128_1	.7021	**.7245**	**.7245**	**.7245**	512_1	.6897	.7066	**.7088**	**.7088**
128_2	.7021	.7422	**.7424**	.7414	512_2	.7135	.7500	**.7520**	.7502
128_4	.7254	.7947	**.7959**	.7921	512_4	.7200	.7819	**.7888**	.7769
128_8	.7142	.7919	**.8050**	.7807	512_8	.7206	.7852	**.7920**	.7837
256_1	.7021	.7217	**.7221**	**.7221**	1024_1	.6969	.7160	**.7174**	**.7174**
256_2	.7066	.7441	**.7448**	.7443	1024_2	.7146	.7509	**.7532**	.7518
256_4	.7235	.7910	**.7940**	.7882	1024_4	.7246	.7840	**.7910**	.7778
256_8	.7166	.7813	**.7987**	.7815	1024_8	.7216	.7835	**.7893**	.7817

Table 2. Average efficiency of $LS_{\bar{f}}$ on NK landscapes, without (0) or after 200 perturbations. We also report comparison with local search (LS for hill-climbing, ILS for iterated local search) on original NK landscapes. The ILS column contains the best average results obtained among the 11 tested parameterizations.

Acknowledgment. This work was partially supported by the Fondation mathématique Jacques Hadamard within the Gaspard Monge Program for Optimization and operations research.

References

1. Matthieu Basseur and Adrien Goëffon. Hill-climbing strategies on various landscapes: an empirical comparison. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference (GECCO)*, pages 479–486. ACM, 2013.
2. 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.
3. Scott Kirkpatrick, C. Daniel Gelatt, and Mario P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
4. Tom Smith, Phil Husbands, Paul J. Layzell, and Michael O'Shea. Fitness landscapes and evolvability. *Evolutionary Computation*, 10:1–34, 2002.