



HAL
open science

On the structure of multiobjective combinatorial search space: MNK-landscapes with correlated objectives

Sébastien Verel, Arnaud Liefoghe, Laetitia Jourdan, Clarisse Dhaenens

► To cite this version:

Sébastien Verel, Arnaud Liefoghe, Laetitia Jourdan, Clarisse Dhaenens. On the structure of multi-objective combinatorial search space: MNK-landscapes with correlated objectives. *European Journal of Operational Research*, 2013, 227 (2), pp.331-342. 10.1016/j.ejor.2012.12.019 . hal-00760097

HAL Id: hal-00760097

<https://hal.science/hal-00760097>

Submitted on 9 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the structure of multiobjective combinatorial search space: *MNK*-landscapes with correlated objectives

Sébastien Verel^{a,c}, Arnaud Liefooghe^{b,c}, Laetitia Jourdan^c, Clarisse Dhaenens^{b,c}

^aUniversité Nice Sophia Antipolis, I3S – CNRS, 06903 Sophia Antipolis, France

^bUniversité Lille 1, LIFL – CNRS, 59655 Villeneuve d'Ascq cedex, France

^cINRIA Lille-Nord Europe, 40 avenue Halley, 59650 Villeneuve d'Ascq, France

Abstract

The structure of the search space explains the behavior of multiobjective search algorithms, and helps to design well-performing approaches. In this work, we analyze the properties of multiobjective combinatorial search spaces, and we pay a particular attention to the correlation between the objective functions. To do so, we extend the multiobjective *NK*-landscapes in order to take the objective correlation into account. We study the co-influence of the problem dimension, the degree of non-linearity, the number of objectives, and the objective correlation on the structure of the Pareto optimal set, in terms of cardinality and number of supported solutions, as well as on the number of Pareto local optima. This work concludes with guidelines for the design of multiobjective local search algorithms, based on the main fitness landscape features.

Key words: multiple objective programming, combinatorial optimization, local search, fitness landscape analysis, Pareto optimal set, Pareto local optima

1. Introduction

In this paper, we study the structure of multiobjective combinatorial search spaces. The aim is to adapt and to enhance the comprehensive design of multiobjective local search approaches, motivated by an *a priori* analysis of problem properties. The attempt of this work is to reduce the gap between the fields of multiobjective combinatorial optimization and of fitness landscape analysis. On the one hand, *multiobjective combinatorial optimization* (MoCO) is one of the most challenging area from multicriteria decision making. Contrary to the single-objective case, there does not exist a single optimal solution, but a set of solutions forming the *Pareto optimal set*. A partial order is defined among feasible solutions, based on the so-called *Pareto dominance relation*. A fundamental issue is related to the identification of this set, or an approximation of it for large-size and difficult MoCO problems. The decision maker then has to select his/her most preferred solution among this set. On the other hand, *fitness landscape analysis* aims to understand the geometry of a combinatorial optimization problem in order to design more efficient search algorithms. However, there is a very little knowledge on the landscape of MoCO problems, where an additional difficulty relies on the structure of the Pareto optimal set. The impact of the main problem-related properties on the behavior and the performance of multiobjective approaches is still far from being well understood.

We here try to give a first step towards understanding the structural properties of MoCO problems.

1.1. Motivations

There exists two main classes of approximate multiobjective search algorithms (Ehrgott and Gandibleux, 2004; Paquete and Stützle, 2007). Their individual dynamics is directly related to the structure of the problem under consideration, itself affected by the problem properties. First, scalar approaches are based on multiple scalarized aggregations of the objective functions, but they are only able to find a subset of Pareto optimal solutions, known as *supported solutions*. Second, Pareto-based approaches are based on the Pareto dominance relation. However, when the size of the Pareto optimal set gets too large, an algorithm should manipulate a limited-size solution set only.

In single-objective optimization, one of the main feature of fitness landscape analysis is related to the number of local optima, to their distribution over the search space and to the shape of their basins of attraction. It has become clear that local optima have a strong impact on the performance of search algorithms. For instance, it has been shown that local optima tend to be clustered in a ‘central massif’ for numerous combinatorial optimization problems, including the family of *NK*-landscapes (Kauffman, 1993). Like in single-objective optimization, local optima clearly have a strong impact on the landscape of a MoCO problem, and then on the behavior and the efficiency of multiobjective approaches. Pareto-based local search algorithms are designed in order to take them into account (Paquete et al., 2007). In general, the aim of such approaches is to find a set of mutually non-dominated local optima. Surprisingly, up to now, there is a lack of study on the

Email addresses: verel@i3s.unice.fr (Sébastien Verel),
arnaud.liefooghe@univ-lille1.fr (Arnaud Liefooghe),
laetitia.jourdan@inria.fr (Laetitia Jourdan),
clarisse.dhaenens@lifl.fr (Clarisse Dhaenens)

number and on the distribution of local optima in MoCO. In this work, one of our attempt is to analyze the structure of the search space, based on a notion of local optimum. However, local optima do not explain all the properties related to MoCO problems. The Pareto optimal set structure by itself also affects the behavior of multiobjective approaches. For the design of multiobjective local search algorithms, the following questions, related to the Pareto optimal set structure and to the number of local optima, are of our interest in this paper:

- (i) What is the cardinality of the Pareto optimal set? Can we intend to identify or approximate the whole set of Pareto optimal solutions, or should we consider a mechanism to bound the size of the approximation set?
- (ii) How many Pareto optimal solutions are supported? Is a scalar approach able to identify or approximate “enough” Pareto optimal solutions?
- (iii) How many local optima, in terms of Pareto dominance, are there in the search space? Can this number be estimated? Is a Pareto-based approach efficient according to such a number of local optima?

A particular attention is paid to the correlation between the objective functions. Indeed, this problem-related property is known to largely affect the solutions of a MoCO problem (Mote and Olson, 1991), and the behavior of search algorithms (Paquete and Stützle, 2006, 2010). To this end, a set of properly-defined problem instances, where all problem-properties can be tuned precisely, is strongly required. Unfortunately, there is a clear lack of *benchmark* MoCO problems. Most existing benchmark instances generally neglect objective correlation, or consist of two-objective problems. To the best of our knowledge, there does not exist any reliable way to set the objective correlation where more than two objective functions are involved. The single exception should be the multiobjective quadratic assignment problem (Knowles and Corne, 2002), where a correlation parameter allows to tune the correlation between different pairs of objective functions, but not for all the objectives at the same time. The remaining properties of problem size, non-linearity and objective space dimension have been recently considered in the family of multiobjective *NK*-landscapes, proposed by Aguirre and Tanaka (2007), but the objective correlation is null.

1.2. Contributions

In this paper, we conduct a fitness landscape analysis for MoCO, based on the multiobjective *NK*-landscapes with objective correlation. We study the structure of a MoCO search space according to the following problem properties: its size, its degree of non-linearity, its objective space dimension (up to five objective functions), and its objective correlation. We measure different features of the search space, related to the cardinality of the Pareto optimal set, the proportion of supported solutions, and the number of local optima. Our results show the importance of taking the objective correlation degree into account on

the design of local search approaches. This paper extends recent works on the multiobjective *NK*-landscapes with objective correlation (Verel et al., 2011a,b). The main contributions of this work can be stated as follows.

- First, we propose a new approach allowing to precisely tune the correlation between multiple objective functions. Benchmark instances are proposed for correlated multi-objective *NK*-landscapes, and we make them available at the following URL: <http://mocobench.sf.net>. Note that our proposal can be generalized to other MoCO problems. We conduct a theoretical analysis and an experimental study to show the sharpness of the objective correlation, between every pair of objective functions.
- Second, we study the co-influence of objective correlation, objective space dimension and non-linearity on the properties of the Pareto optimal set by complete enumeration: its cardinality, the proportion of supported Pareto optimal solutions.
- Third, we show the co-influence of objective correlation, objective space dimension and non-linearity on the number of local optima. We propose a method based on the length of an adaptive walk to estimate this number, and we study the number of local optima for large-size problem instances. Hence, we can give better insights on the structure of the search space for large-size instances.
- At last, the consequences of these properties are discussed for the design of local search algorithms. Guidelines are provided in order to make proper choices for several methodological classes.

The remainder of the paper is organized as follows. Section 2 deals with MoCO, fitness landscapes and multiobjective local search. Section 3 presents a new approach for the design of multiobjective *NK*-landscapes with objective correlation. Section 4 deeply analyzes the Pareto optimal set structure according to objective space dimension, degree of non-linearity, and objective correlation. In Section 5, we study the number of local optima for enumerable instances, we propose a method to estimate this measure, and we put this method in application for large-size instances. At last, the consequences of the co-influence of problem size, objective space dimension, objective correlation and non-linearity on the main properties of the search space are summarized in the last section. This discussion is completed with multiple insights and guidelines on the design of multiobjective local search algorithms, together with open research directions.

2. Local search and fitness landscapes for multiobjective combinatorial optimization

This section introduces definitions for multiobjective combinatorial optimization (MoCO), and discusses fitness landscapes as well as the design of local search algorithms in such a context. Table 1 summarizes the main notations used in the paper.

Table 1: Main notations used in the paper.

Notation	Meaning
\mathcal{X}	set of feasible solutions in the decision space
\mathcal{Z}	set of feasible outcome vectors in the objective space
\mathcal{X}_{PO}	set of Pareto optimal solutions
PLO	a Pareto local optimum solution
\mathcal{X}_{PLO}	set of Pareto local optima
\mathcal{X}_{SPO}	set of supported Pareto optimal solutions

2.1. Multiobjective combinatorial optimization

A MoCO problem can be defined by a set of $M \geq 2$ objective functions $f = (f_1, f_2, \dots, f_M)$, and a (discrete) set \mathcal{X} of feasible solutions in the *decision space*. Let $\mathcal{Z} = f(\mathcal{X}) \subseteq \mathbf{R}^M$ be the set of feasible outcome vectors in the *objective space*. In a maximization context, a solution $x' \in \mathcal{X}$ is dominated by a solution $x \in \mathcal{X}$, denoted by $x' < x$, iff $\forall i \in \{1, 2, \dots, M\}$, $f_i(x') \leq f_i(x)$ and $\exists j \in \{1, 2, \dots, M\}$ such that $f_j(x') < f_j(x)$. A solution $x^* \in \mathcal{X}$ is said to be *Pareto optimal* (or *efficient*, *non-dominated*), if there does not exist any other solution $x \in \mathcal{X}$ such that $x^* < x$. The set of all Pareto optimal solutions is called the *Pareto optimal set* (or the *efficient set*). Its mapping in the objective space is called the *Pareto front*. A possible approach in MoCO is to identify a minimal complete Pareto optimal set, here denoted by \mathcal{X}_{PO} , *i.e.* for each point from the Pareto front, a single Pareto optimal solution in the decision space is considered, even if there could be more than one.

Generating the entire Pareto optimal set is often infeasible for two main reasons. First, for most MoCO problems, the number of Pareto optimal solutions is typically exponential in the size of the problem instance (Ehrgott, 2005). In that sense, most MoCO problems are said to be *intractable*. Second, deciding if a feasible solution belongs to the Pareto optimal set is NP-complete for numerous MoCO problems (Serafini, 1987), even if none of its single-objective counterpart is NP-hard. Therefore, the overall goal is often to identify a good Pareto set approximation. To this end, heuristics in general, and evolutionary algorithms in particular, have received a growing interest since the late eighties. Multiobjective heuristics still constitute an active research area (Deb, 2001; Ehrgott and Gandibleux, 2004).

2.2. Multiobjective fitness landscapes

A *neighborhood operator* is a function $\mathcal{N} : \mathcal{X} \rightarrow 2^{\mathcal{X}}$ that assigns a set of solutions $\mathcal{N}(x) \subset \mathcal{X}$ to any solution $x \in \mathcal{X}$. The set $\mathcal{N}(x)$ is called the *neighborhood* of x , and a solution $x' \in \mathcal{N}(x)$ is called a *neighbor* of x . In single-objective combinatorial optimization, a fitness landscape can be defined by the triplet $(\mathcal{X}, \mathcal{N}, h)$, where $h : \mathcal{X} \rightarrow \mathbf{R}$ represents the fitness function, that can be pictured as the *height* of the corresponding solutions. Each peak of the landscape corresponds to a local optimum. In a single-objective maximization context, a *local optimum* is a solution $x^* \in \mathcal{X}$ such that $\forall x \in \mathcal{N}(x^*)$, $h(x) \leq h(x^*)$. The efficiency of local search algorithms has been shown to be

related to the number of local optima for the problem under study, and to their distribution over the landscape (Merz, 2004).

In MoCO, given that Pareto optimal solutions are to be found, the notion of local optimum has to be defined in terms of Pareto optimality. Let us define the concepts of Pareto local optimum and of Pareto local optimum set. For more details, refer to Paquete et al. (2007). A solution $x \in \mathcal{X}$ is a *Pareto local optimum (PLO)* with respect to a neighborhood structure \mathcal{N} if there does not exist any neighboring solution $x' \in \mathcal{N}(x)$ such that $x < x'$. The *Pareto local optimum set* $\mathcal{X}_{PLO} \in \mathcal{X}$ is the set of whole *PLO*. A *Locally non-dominated set* $A \in \mathcal{X}$ with respect to a neighborhood structure \mathcal{N} is a set of mutually non-dominated solutions such that $\forall x \in A$, $\forall x' \in (\mathcal{N}(x) \setminus A)$, there exists $x'' \in A$ such that $x' < x''$. In other words, a locally non-dominated set cannot be improved, in terms of Pareto optimality, by adding neighboring solutions. As a consequence, a locally non-dominated set is a subset of \mathcal{X}_{PLO} . Let us note that a Pareto optimal solution is a *PLO*, and that \mathcal{X}_{PO} is a locally non-dominated set.

A multi-objective fitness landscape can be defined by the triplet $(\mathcal{X}, \mathcal{N}, f)$, where $f : \mathcal{X} \rightarrow \mathbf{R}^M$ represents the multi-dimensional objective function. There exists a small amount of literature related to fitness landscape for MoCO. Borges and Hansen (1998) study the distribution of local optima, in terms of scalarized functions, for the multiobjective traveling salesman problem (TSP). Another analysis of neighborhood-related properties for biobjective TSP instances of different structures is given by Paquete and Stützle (2009). Knowles and Corne (2002) lead a landscape analysis on the multiobjective quadratic assignment problem with a rough objective correlation. Garrett and Dasgupta (2007) discuss standard tools from fitness landscape analysis to MoCO, and an experimental study is conducted with fitness distance correlation. But this measure requires the true Pareto optimal set to be known. Afterwards, Garrett and Dasgupta (2009) consider the fitness landscape of a MoCO problem as a neutral landscape, where many solutions are incomparable. At last, in previous works on multiobjective *NK* landscapes (Aguirre and Tanaka, 2007), enumerable fitness landscapes are divided into different fronts, *i.e.* layers of mutually non-dominated solutions, following the dominance principles of NSGA-II (Deb, 2001). These landscapes are studied according to the number of fronts, the number of solutions on each front, the probability to move from one front to another, and the hypervolume of the Pareto front. However, this study allows to analyze only small search spaces from the single point of view of dominance rank.

2.3. Multiobjective local search

Initial approaches dealing with MoCO were based on successive transformations of the original multiobjective problem into single-objective ones by means of a scalarization strategy. Scalar approaches are generally based on a weighted-sum aggregation of the objective functions, that can be defined as follows. $\forall x \in \mathcal{X}$:

$$f_{\lambda}(x) = \sum_{i=1}^M \lambda_i f_i(x) \quad (1)$$

where $\lambda_i > 0$ for all $i \in \{1, \dots, M\}$. The problem is now to identify a (single) solution that maximizes f_λ . For any given weighting coefficient vector λ , if $x^* = \arg \max_{x \in X} f_\lambda(x)$, then x^* is a supported Pareto optimal solution (Ehrgott, 2005). Multiple weighting coefficient vectors can be iteratively defined so that several non-dominated solutions are identified (or approximated). For each scalarization, the corresponding solution is incorporated into an approximation set, whose dominated solutions are then discarded. However, in the combinatorial case, a number of Pareto optimal solutions, known as non-supported solutions, are not optimal for any definition of f_λ . A solution $x \in X$ is a *non-supported Pareto optimal solution* if (i) $x \in X_{PO}$ (i.e. x is a Pareto optimal solution), (ii) $\forall \lambda \in \mathbf{R}^M$ such that $\lambda_i > 0, i \in \{1, \dots, M\}, x \neq \arg \max_{x' \in X} f_\lambda(x')$. On the contrary, there exist *supported (Pareto optimal) solutions*, that are solutions whose corresponding objective vectors are located on the convex hull of the Pareto front. The set of all supported Pareto optimal solutions will be denoted by X_{SPO} . As a consequence, the proportion of non-supported solutions over the Pareto optimal set has a direct impact on the ability of scalar approaches to find a proper Pareto set approximation.

Over the years, other types of approaches were proposed. Their internal search mechanisms are based on the explicit or implicit use of the Pareto dominance relation, that allows to define a partial order between feasible solutions. The basic idea is to maintain a set of solutions (typically an archive of mutually non-dominated solutions). The content of this set is then iteratively updated with new solutions built by means of neighborhood operators. The update of this set is based on a decision that specifies which solutions to accept or to choose for the next iteration. This process is iterated until no further improvement is possible or another stopping condition is fulfilled. In the end, this set corresponds to the approximation outputted by the algorithm. A whole family of such local search algorithms follows this general scheme (Liefoghe et al., 2012). One of them is the Pareto Local Search (PLS), proposed by Paquete et al. (2007). PLS combines the use of a neighborhood structure with the management of an archive of mutually non-dominated solutions found so far. The basic idea is to iteratively improve this archive by exploring the neighborhood of its own content until no further improvement is possible, i.e. the archive falls in a Pareto local optimum set. Hence, the behavior of a PLS-like algorithm clearly depends on the number and the distribution of PLO found along the search process. Indeed, when an approximation set contains no PLO, improving solutions can be identified by exploring the neighborhood of solutions from the set. On the contrary, when an approximation set contains a large proportion of PLO, it becomes harder to find additional non-dominated solutions. In general, the implicit goal of Pareto-based approaches is to identify an approximation whose image in the objective space is (i) close to and (ii) well-spread along the Pareto front. However, as the number of Pareto optimal solutions is generally intractable, we often have to design specific strategies to limit the size of the approximation set for large-size problem instances (Knowles and Corne, 2004). As a consequence, the *cardinality* of the Pareto optimal set also plays a major role on the design of multiobjective local search algo-

rithms.

3. Multiobjective NK -landscapes with objective correlation

In single-objective optimization, the family of NK -landscapes forms an interesting model to study the influence of non-linearity (epistasis) on the number of local optima (Kauffman, 1993). In this section, we define the class of ρMNK -landscapes, which extend the multiobjective NK -landscapes recently proposed by Aguirre and Tanaka (2007). In our multiobjective model, the correlation between objective functions can be precisely tuned by a correlation coefficient value. It allows to study the simultaneous influence of problem size, non-linearity, objective space dimension and objective correlation on the main properties of multiobjective fitness landscapes. We first introduce existing single-objective and multiobjective NK -landscapes. Then, the construction of ρMNK -landscapes is defined and the analytic proof of the correlation between objectives, completed with an experimental study, are given.

3.1. NK - and MNK -landscapes

The family of NK -landscapes is a problem-independent model used for constructing multimodal landscapes (Kauffman, 1993). N refers to the number of bits in the decision space (i.e. the bit string length) and K to the number of bits that influence a particular bit from the string (the epistatic interactions or degree). By increasing the value of K from 0 to $(N - 1)$, NK -landscapes can be gradually tuned from smooth to rugged. The set of feasible solutions is made of binary strings $\mathcal{X} = \{0, 1\}^N$. The fitness function (to be maximized) $f_{NK} : \mathcal{X} \rightarrow [0, 1)$ associates to each solution a real number between 0 and 1. It can be computed as follows: An ‘atom’ with a fixed epistasis level is represented by a fitness component $f_i : \{0, 1\}^{K+1} \rightarrow [0, 1)$ associated to each bit $i \in \{1, \dots, N\}$. Its value depends on the value of bit i and also on the values of K other bit positions (K must fall between 0 and $N - 1$). In other words, there are K bits interacting with bit i . The fitness $f_{NK}(x)$ of a solution $x \in \mathcal{X}$ corresponds to the mean value of its N fitness components f_i :

$$f_{NK}(x) = \frac{1}{N} \sum_{i=1}^N f_i(x_i, x_{i_1}, \dots, x_{i_K}) \quad (2)$$

where $\{i_1, \dots, i_K\} \subset \{1, \dots, i - 1, i + 1, \dots, N\}$, and x_i is the value of bit i in solution x . According to Eq. (2), the parameter K tunes the degree of non-linearity (epistasis). Several ways have been proposed to set the K bits from the bit strings. Two possibilities are mainly used: adjacent and random positions. With an adjacent position, the nearest K bits to bit $i \in N$ are chosen (the bit string is chosen to have periodic boundaries). In this work, we set the K bits randomly on the bit string of size N in order to avoid the exploitation of adjacent positions in order to improve the performance of algorithms (Weinberger, 1996). Indeed, such particular interactions can introduce some bias in the structure of the problem, which is beyond the scope of our study.

Each fitness component f_i is specified by extension, i.e. a number $y_{x_i, x_{i_1}, \dots, x_{i_K}}^i$ from $[0, 1)$ is associated with each element

$(x_i, x_{i_1}, \dots, x_{i_k})$ from $\{0, 1\}^{K+1}$. Those numbers are uniformly distributed in the range $[0, 1)$. As a consequence, it is very unlikely that the same fitness value is assigned to two different solutions.

More recently, a multiobjective variant of NK -landscapes (namely MNK -landscapes) has been defined with a set of M fitness functions (Aguirre and Tanaka, 2007). $\forall m \in \{1, \dots, M\}$ $\forall x \in \mathcal{X}$:

$$f_{NK_m}^m(x) = \frac{1}{N} \sum_{i=1}^N f_i^m(x_i, x_{i_1}^m, \dots, x_{i_{k_m}}^m) \quad (3)$$

The number of interactions between bits K_m (also known as *epistasis degree*) can theoretically be different for each fitness function $f_{NK_m}^m$ with $m \in \{1, \dots, M\}$. But, in their analysis, Aguirre and Tanaka (2007) only looked at landscapes with the same epistasis degree $K_m = K$ for all the objective functions. Similarly, we here analyze landscapes with the same epistasis degree $K_m = K$ for all the objective functions. Each fitness component f_i^m is specified by extension with the numbers $y_{x_i, x_{i_1}^m, \dots, x_{i_{k_m}}^m}^{m,i}$. In the original MNK -landscapes, these numbers are randomly and independently drawn from $[0, 1)$. We here propose an approach to design MNK -landscapes with correlated objective functions.

3.2. ρMNK -landscapes

First, let us define $CMNK$ -landscapes, where the epistasis structure is identical for all the objectives: $\forall m \in \{1, \dots, M\}$, $K_m = K$ and $\forall m \in \{1, \dots, M\}$, $\forall j \in \{1, \dots, K\}$, $i_j^m = i_j$. The fitness components are not defined independently. Indeed, the M numbers $(y_{x_i, x_{i_1}, \dots, x_{i_k}}^{1,i}, \dots, y_{x_i, x_{i_1}, \dots, x_{i_k}}^{M,i})$, associated with each element $(x_i, x_{i_1}, \dots, x_{i_k})$, follow a multivariate uniform distribution of dimension M , defined by a correlation matrix C . Thus, the y 's follow a multidimensional distribution with uniform marginals and the correlations between $y^{m,i}$'s are defined by the matrix C .

The matrix C is a correlation matrix. As a consequence, it is a symmetric positive-definite matrix where $\frac{M(M-1)}{2}$ numbers can be defined. In order to limit the number of parameters, that tune the problem instance, we define the matrix $C_\rho = (c_{np})$ with the same correlation between all the objective functions: $c_{mm} = 1$ for all n , and $c_{np} = \rho$ for all $n \neq p$. In this case, we denote $CMNK$ -landscapes by ρMNK -landscapes. Let us note that, when $\rho = 0$, our model is not exactly equivalent to the classical MNK -landscapes proposed by Aguirre and Tanaka (2007). Indeed fitness values are drawn independently, but the two models are different in the epistasis structure.

For obvious reasons, it is not possible, for any MoCO problem, to have a high negative correlation between all pairs of objective functions. Then, for the ρMNK -landscapes, we cannot construct the matrix C_ρ for all ρ values between $[-1, 1]$. Indeed, C_ρ must be positive-definite: $\forall u \in \mathbb{R}^M$, $u^t C_\rho u \geq 0$. So, ρ must be greater than $\frac{-1}{M-1}$. For two-objective problems, all the correlation coefficients in $[-1, 1]$ are possible. However, for three-objective problems, the correlation ρ must fall in $[-0.5, 1]$. Of course, if one wants to study very negative correlations between some pairs of objectives, it is possible to design a

matrix C that holds the condition that C is positive-definite. To generate random variables with uniform marginals and a specified correlation matrix C , we follow the work of Hotelling and Pabst (1936). We first generate a (Z_1, \dots, Z_M) vector following a multi-normal distribution of means 0 and correlation matrix $R = 2 \sin(\frac{\pi}{6}C)$. Then, the values $z_i = \Phi(Z_i)$ are uniformly distributed with a correlation matrix C , where Φ is the univariate normal cumulative density function. Note that this is not the only way to generate a multivariate uniform distribution.

3.3. Objective correlation

The objective correlation is the correlation between solutions in the objective space. The construction of $CMNK$ -landscapes defines correlations between the y 's but not directly between the objectives. In this section, we prove by algebra that the correlation between objectives is exactly tuned by the matrix C . This proof is completed by a short experimental study.

3.3.1. Theoretical analysis

This section gives the relation between the matrix C which defines the $CMNK$ -landscapes and the objective correlation.

Proposition 1. *Let $C = (c_{np}) \in \mathbb{R}^M \times \mathbb{R}^M$ be a correlation matrix. On the class of $CMNK$ -landscapes, the expected value of the correlation between the objective functions n and p is given by the coefficient c_{np} of the matrix C .*

Proof. Let $F^m = (f_{NK}^m(x))$ be the fitness vector values of the 2^N solutions with respect to objective m . The correlation between the objectives n and p is: $cor(F^n, F^p) = \frac{cov(F^n, F^p)}{\sigma_n \sigma_p}$ where σ_n and σ_p are the standard deviations of fitness values over the landscape of the n^{th} and p^{th} objective functions. F^n (resp. F^p) corresponds to the average value of the N vectors F_i^n (resp. F_i^p) of fitness component values:

$$cov(F^n, F^p) = \frac{1}{N^2} \sum_{i,j=1}^N cov(F_i^n, F_j^p) \quad (4)$$

By definition, when $i \neq j$, $cov(F_i^n, F_j^p) = 0$ and $cov(F_i^n, F_i^p) = c_{np} \cdot \sigma_{ni} \cdot \sigma_{pi}$, where c_{np} is the correlation defined in the matrix C , and σ_{ni} (resp. σ_{pi}) is the standard deviation of fitness component i . The correlation between objectives n and p becomes:

$$cor(F^n, F^p) = c_{np} \frac{\sum_{i=1}^N \sigma_{ni} \sigma_{pi}}{N^2 \sigma_n \sigma_p} \quad (5)$$

By construction of the fitness functions, the following relation between standard deviations stands $\sigma_n^2 = \frac{1}{N} \sum_{i=1}^N \sigma_{ni}^2$ (resp. for σ_p^2). On average, the σ_{ni} are equal to the standard deviation of the uniform distribution on $[0, 1)$.

$$E(cov(F^n, F^p)) = c_{np} \quad (6)$$

□

Then, the expected value of the correlations between objective functions are given by the matrix C . In the ρMNK -landscapes, the parameter ρ allows to tune very precisely the correlation between all pairs of objectives.

3.3.2. Experimental study

In order to validate the behavior of the correlation coefficient given by Eq. (6) experimentally, we conduct an empirical study for $N = 18$ in order to enumerate the search space exhaustively. To minimize the influence of the random creation of landscapes, we considered 30 different and independent landscapes for each parameter combinations: ρ , M , and K . Measures reported are the average over these 30 landscapes. The remaining set of parameters are given in Table 2. Figure 1 gives the average value¹ of the Spearman correlation coefficient according to parameters ρ , M and K . This confirms the result of Eq. (6). Within 30 instances, the correlation coefficient is very close to the expected value ρ between all pairs of objective functions.

3.4. Discussion

To summarize, the four parameters of the family of ρMNK -landscapes are: (i) the length of bit strings N , (ii) the number of epistatic links K , (iii) the number of objective functions M , and (iv) the correlation coefficient ρ . The parameter ρ allows to tune the objective correlation very precisely. Note that the approach proposed to tune the objective correlation can also be applied to other MoCO problems where the objective functions are summing objectives, share the same definition, but are computed with different cost or profit matrices. This is the case, for instance, of most MoCO problems presented by Ehrgott (2005), including the multiobjective versions of the shortest path problem, minimum spanning tree problem, knapsack problem and traveling salesman problem, among others. Our approach can also be generalized to the previous attempt made in that direction for the multiobjective quadratic assignment problem (Knowles and Corne, 2002). However, here the objective correlation can be tuned between every pair of objective functions, and note only between one arbitrary objective function with each remaining one, which is not reliable for problems with more than two objective functions. The source code of the instance generator and the set of benchmark instances used in the paper can be retrieved at the following URL: <http://mocobench.sf.net>.

In the following, we study the influence of the problem dimension, the non-linearity (epistasis), the number of objective functions and the objective correlation on the properties of the Pareto optimal set and of Pareto local optima for the ρMNK -landscapes.

4. Analysis of the Pareto optimal set

In this section, we conduct an experimental analysis on the ρMNK -landscapes in order to study different properties of the Pareto optimal set (X_{PO}): its cardinality and the proportion of extreme supported solutions. We conduct an empirical study for $N = 18$ so that we can identify the (exact) Pareto optimal set by a complete enumeration of the decision space. Measures reported are the average values over 30 different and independent

¹For $M > 2$, there are several correlation coefficients. We report here the average correlation coefficients over all the objectives.

landscapes. The parameters under investigation in this study are given in Table 2.

4.1. Cardinality of the Pareto optimal set

Figure 2 shows the proportion of Pareto optimal solutions in the search space according to parameters K , ρ and M of ρMNK -landscapes. First of all, the epistatic parameter K has a low influence on the results. This is in accordance with the results reported by Aguirre and Tanaka (2007), where it was shown that the number of epistatic interactions is a weaker factor compared to the number of objectives. At the opposite, the objective correlation coefficient ρ modifies the number of Pareto optimal solutions to several orders of magnitude. Indeed, the proportion decreases from 10^{-4} for $\rho = -0.9$ to 10^{-5} for $\rho = 0.9$ for two-objective problems, and from 10^{-1} for $\rho = -0.2$ to 10^{-5} for $\rho = 0.9$ for five-objective problems. With respect to the number of objective functions ($M = 2, 3$, and 5), the size increases of several orders of magnitude with M . For a negative objective correlation ($\rho = -0.2$), the proportion of Pareto optimal solutions goes from 10^{-4} up to 10^{-1} , whereas it goes from 10^{-5} up to 10^{-4} for a positive correlation ($\rho = 0.9$).

The influence of objective correlation on the size of X_{PO} is at least as important as the objective space dimension. A lot of solutions are Pareto optimal when the correlation is highly negative. Now, let us suppose that we want to set, or to bound, the size of an approximation set by 100 solutions, *i.e.* a proportion of $4 \cdot 10^{-4}$ solutions with respect to the search space size for $N = 18$. Such a parameter setting is often used while handling an archive of non-dominated solutions in any search approach. For any correlation value ρ , a 100-solution approximation set always allows to store the whole Pareto optimal set for two-objective problems. However, this is not the case for a higher objective space dimension. For instance, for $M = 5$, 100 solutions suffice to store the whole Pareto optimal set for a high objective correlation only ($\rho > 0.5$). In other words, for $\rho < 0.5$, we cannot intend to identify all the Pareto optimal solutions exhaustively by handling a 100-solution approximation set.

To summarize, when the number of objectives increases, and even more when the objective correlation becomes negative, the size of the Pareto optimal set is very large, and then quickly becomes intractable. In such a case, it is not reasonable to identify the whole Pareto optimal set, and a limited-size approximation should be considered. This first result shows the importance to design a benchmark where the objective correlation can be tuned precisely, even when $M > 2$. Such a property should be taken into consideration for the development of local search methods, when the number of objectives becomes too large, and when there is a very negative correlation between the objective functions. A special attention should be paid to the bounding-size mechanism of the approximation set maintained by the algorithm.

4.2. Proportion of supported Pareto optimal solutions

The number of extreme supported solutions in the search space roughly follows the number of Pareto optimal solutions.

Table 2: Parameter setting used in the paper for the experimental analysis (30 random landscapes are created for each parameter combination: ρ , M , N and K).

Description	Parameter	Set of values
problem size	N	(see text for details)
epistasis	K	{2, 4, 6, 8, 10}
number of objectives	M	{2, 3, 5}
objective correlation	ρ	{-0.9, -0.7, -0.4, -0.2, 0.0, 0.2, 0.4, 0.7, 0.9} such that $\rho \geq \frac{-1}{M-1}$

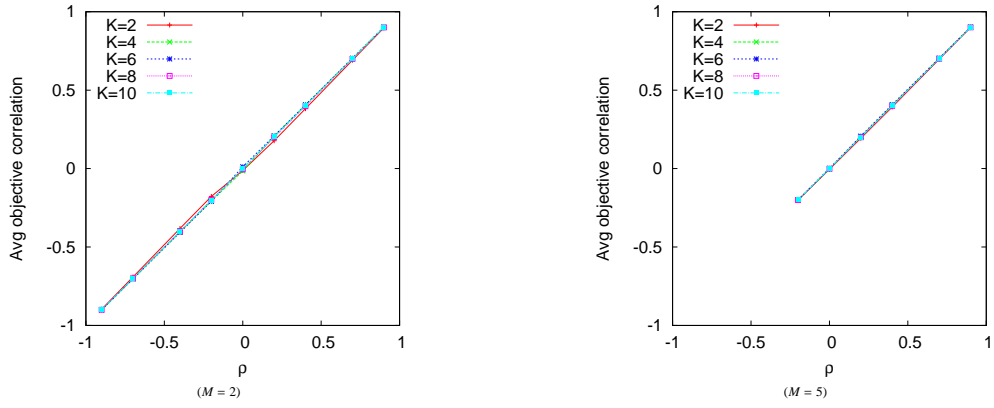


Figure 1: Average values of the objective correlation according to the parameter ρ . The number of objectives is $M = 2$ (left) and $M = 5$ (right).

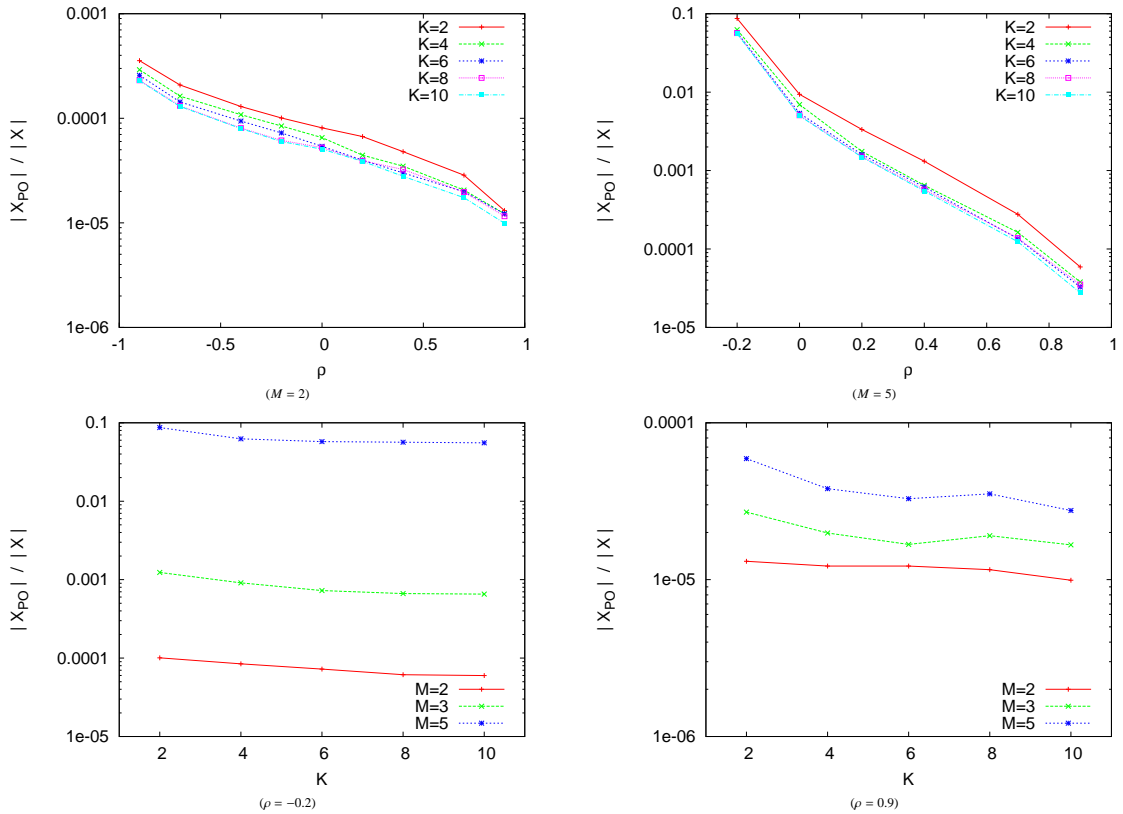


Figure 2: Average ratio of the number of Pareto optimal solutions compared to the size of the search space (2^N) according to parameter ρ for different K -values (top left $M = 2$, right $M = 5$), and according to parameter K for different M -values (bottom left $\rho = -0.2$, right $\rho = 0.9$). Notice the log y-scale and the different y-ranges.

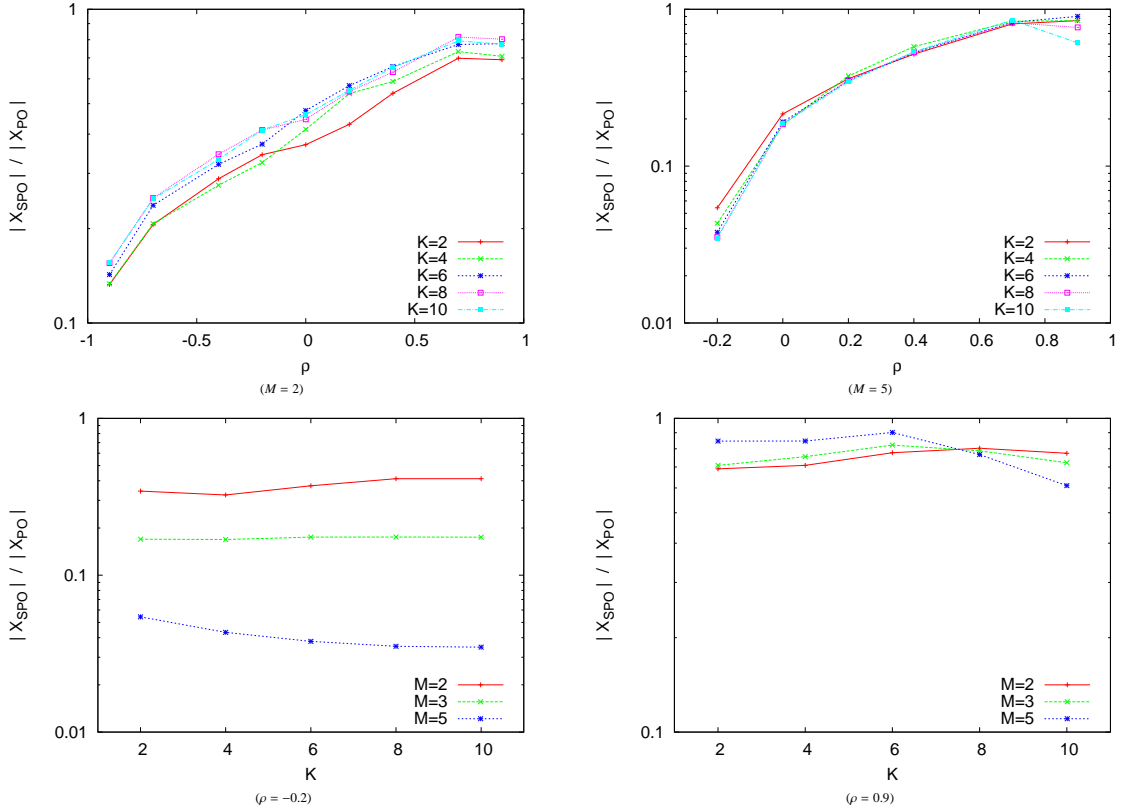


Figure 3: Average ratio of the number of supported Pareto optimal solutions compared to the size of the Pareto optimal set according to parameter ρ for different K -values (top left $M = 2$, right $M = 5$), and according to parameter K for different M -values (bottom left $\rho = -0.2$, right $\rho = 0.9$). Notice the log y-scale.

The epistatic parameter K has low influence on the size. When the objective space dimension increases, or the objective correlation decreases, the number of supported solution gets higher. The proportion of supported solutions in the Pareto optimal set is given in Figure 3. This proportion is nearly independent of the epistasis degree (K). But note that it may change with a landscape using a different epistasis structure on every objective function. However, when the objective correlation increases, this proportion increases. For high objective correlation ($\rho = 0.9$), nearly all solutions become supported (this is even the case for some instances). The same observation can be made with respect to the number of objectives: the proportion of supported solutions increases with the cardinality of the Pareto optimal set, but the former increases faster than the latter.

While putting this property in relation with the design of local search algorithms, we can conclude that scalar approaches should become more appropriate when the number of objectives is low, and when the objective correlation is high. Indeed, in such a case, most Pareto optimal solutions are supported, so that they correspond to optimal solutions for a weighted-sum aggregation problem and a given setting of the weighting coefficient vector.

4.3. Discussion

In this section, we analyzed the consequence of the objective space dimension, the non-linearity, and the objective correla-

tion on the structure of the Pareto optimal set. Figure 4 shows three examples of ρMNK -landscapes represented in the objective space. The number of objectives is $M = 2$, the epistasis parameter is $K = 4$ and the bit string length is $N = 18$. The figure gives a summary of the results in a more intuitive way. When the objective correlation is negative, the objective functions are in conflict with each other (feasible solutions are in green). The Pareto optimal set (in red) is large. In such a case, a local search has to find a limited-size Pareto set approximation only, and not the whole Pareto optimal set. When the objective correlation is null, the image of the search space in the objective space can be represented as a multidimensional “ball”. The objective functions are independent, as previous studies based on the MNK -landscapes (Aguirre and Tanaka, 2007). When the objective correlation is positive, there exist few solutions in the Pareto optimal set. Nearly all Pareto optimal solutions become supported. Indeed, when the number of objectives is low, and when the objective correlation is high, Pareto optimal solutions are supported. We can conclude that scalar approaches should become more appropriate.

Bringing those properties with the design of local search algorithms helps to take proper decisions for several methodological issues. Moreover, while keeping the intensification vs. diversification trade-off in mind, it seems that neighboring or variation operators promoting exploitation should be more appropriate when the objective correlation is high. On the con-

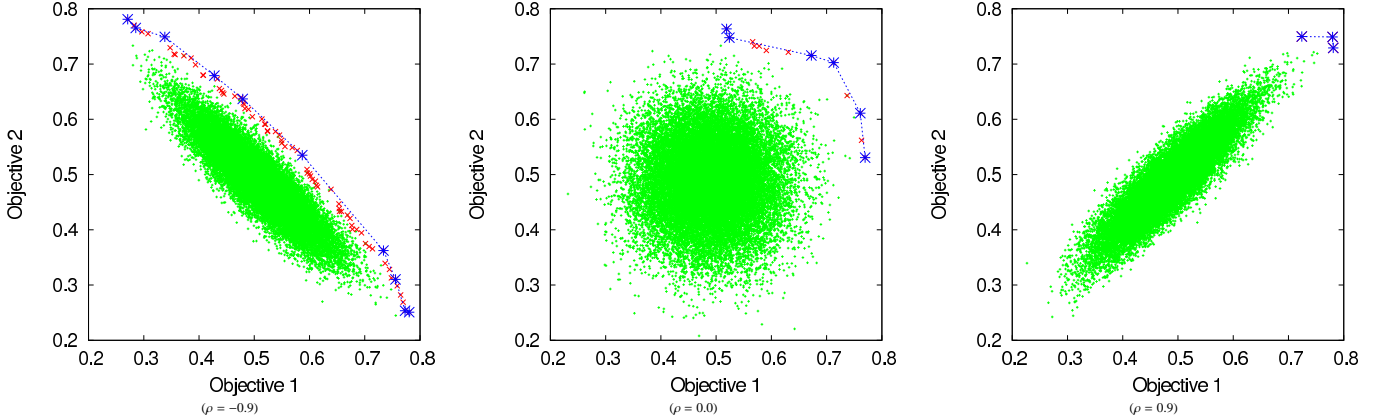


Figure 4: The objective space (maximization problem) for three landscapes. The number of objectives is $M = 2$, the length of bit strings is $N = 18$, the epistasis parameter is $K = 4$. From left to right, the correlation increases from negative correlation to positive correlation ($\rho = -0.9, 0.0$ and 0.9). Green points are random solutions from the search space (10% of the size), red points are Pareto optimal solutions, and blue points are supported Pareto optimal solutions.

rary, exploration has more importance for conflicting objectives, since many Pareto optimal solutions are to be found. However, the Pareto optimal set does not cover all the search space properties. In the next section, we focus on the properties related to Pareto local optima.

5. Analysis of Pareto local optima

In this section, we first study the number of Pareto local optima according to the objective correlation, the number of objectives and the epistasis of ρMNK -landscapes. Next, we analyze its relation with the size of the Pareto optimal set. Then, we propose an adaptive walk that is able to estimate the number of Pareto local optima. At last, we estimate the number of Pareto local optima for large-size problems. The instances under study are defined by the parameter setting given in Table 2.

5.1. Number of Pareto local optima

First, let us remind that a Pareto Local Optimum (PLO) is a solution with no neighbor that dominates it (see Section 2.2). Figure 5 shows the average number of PLO to the size of the search space for different ρMNK -landscapes parameter settings. PLO are enumerated exhaustively for $N = 18$. As the well-known result from single-objective NK -landscapes (Kauffman, 1993), the number of PLO increases with the epistasis degree. For instance, with an objective space dimension $M = 2$ and an objective correlation $\rho = 0.9$, the average number of PLO increases more than 30 times: from 192 for $K = 2$ to 6048 for $K = 10$. However, the range of PLO is larger with respect to objective correlation. For the same epistatic degree and number of objectives, the number of PLO decreases exponentially (Figure 5, top). Indeed, for an objective space dimension $M = 2$ and an epistasis degree $K = 4$, the average number of PLO decreases more than 120 times: from 82,093 for a negative correlation ($\rho = -0.9$) to 672 for a positive correlation ($\rho = 0.9$).

This result can be interpreted as follows. Let us consider an arbitrary solution $x \in \mathcal{X}$, and two different objective functions

f_n, f_p such that $n \neq p$. When the objective correlation is highly positive, there is a high probability for $f_n(x)$ to be close to $f_p(x)$. In the same way, the fitness values $f_n(x')$ and $f_p(x')$ of a given neighbor $x' \in \mathcal{N}(x)$ are probably close. So, for a given solution x such that it exists a neighbor $x' \in \mathcal{N}(x)$ with a better f_n -value, the probability is high that $f_p(x')$ is better than $f_p(x)$. More formally, the probability $\mathbf{P}(f_p(x') > f_p(x) \mid f_n(x') > f_n(x))$, with $x \in \mathcal{X}$ and $x' \in \mathcal{N}(x)$, increases with the objective correlation. Then, a solution has a higher probability of being dominated when the objective correlation is high. Moreover, the probability that a solution dominates all its neighbors decreases with the number of objectives. Figure 5 (bottom) corroborates this hypothesis. When the objective correlation is negative ($\rho = -0.2$), the number of PLO changes in an order of magnitude from $M = 2$ to $M = 3$, and from $M = 3$ to $M = 5$. This range is smaller when the correlation is positive. When the number of objectives is large and the objective correlation is negative, almost all solutions are PLO.

Assuming that the difficulty for Pareto-based search approaches gets higher when the number of PLO is large, the difficulty of ρMNK -landscapes increases when: (i) the epistasis increases, (ii) the number of objective functions increases, (iii) the objective correlation is negative, and its absolute value increases. Section 5.4 will precise the relative difficulty related to those parameters for large-size problem instances.

5.2. Estimating the cardinality of the Pareto optimal set?

When the number of Pareto optimal solutions is too large, it becomes impossible to enumerate them all. A search approach should then manipulate a limited-size solution set only during the search. The cardinality of the Pareto optimal set has been shown to play a major role in the design of multiobjective local search in Section 4.1. It would then be convenient to approximate the size of the Pareto optimal set from the number of PLO. Figure 6 shows the scatter plot of the average size of the Pareto optimal set vs. the average number of PLO in log-scales. Points are scattered over the regression line with the Spearson correlation coefficient of 0.82, and the regression line equation

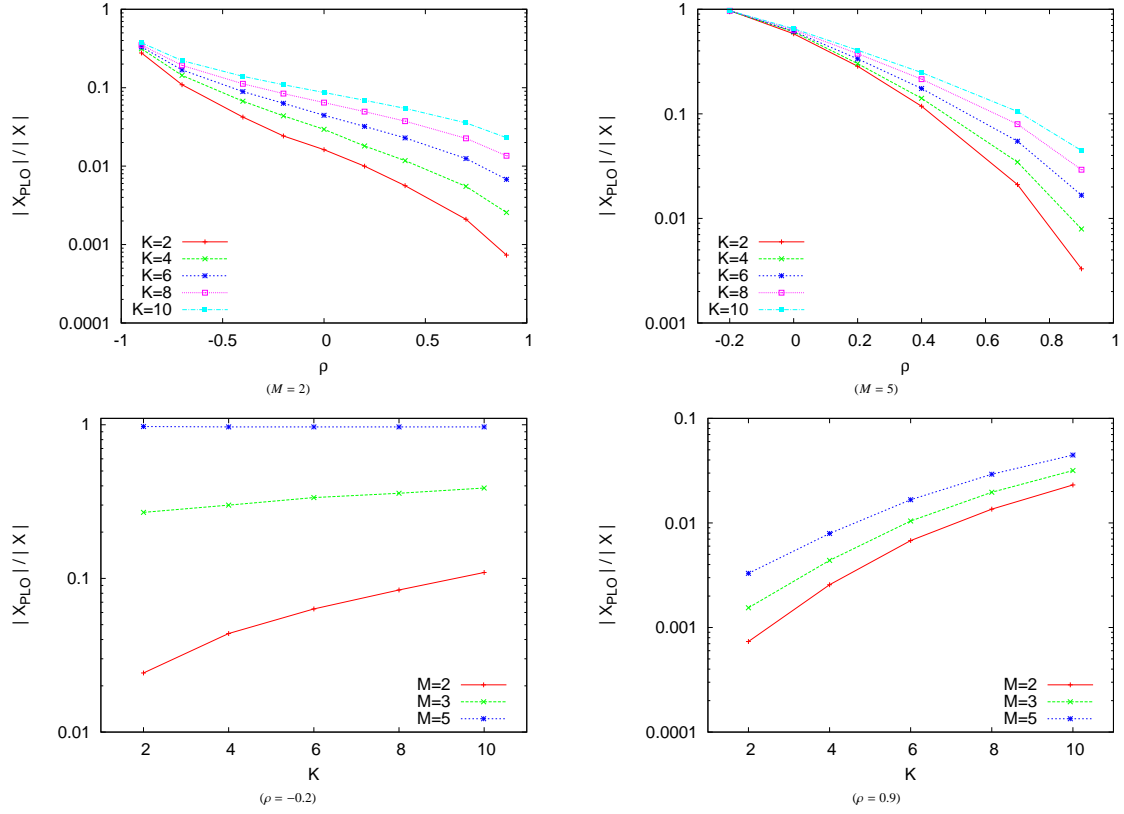


Figure 5: Average number of PLO to the size of the search space according to parameter ρ for different K -values (top left $M = 2$, right $M = 5$), and to parameter K for different M -values (bottom left $\rho = -0.2$, right $\rho = 0.9$). The problem size is $N = 18$.

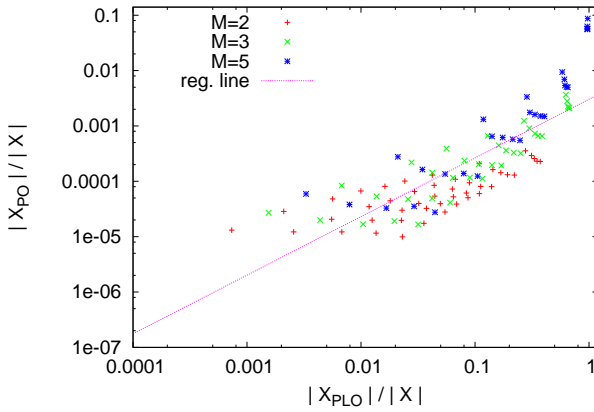


Figure 6: Scatter plot of the average size of the Pareto optimal set (to the size of the search space) vs. the average number of PLO (to the size of the search space) for the 110 possible combinations of parameters. The problem size is $N = 18$. The correlation coefficient is 0.82. Notice the log-scales.

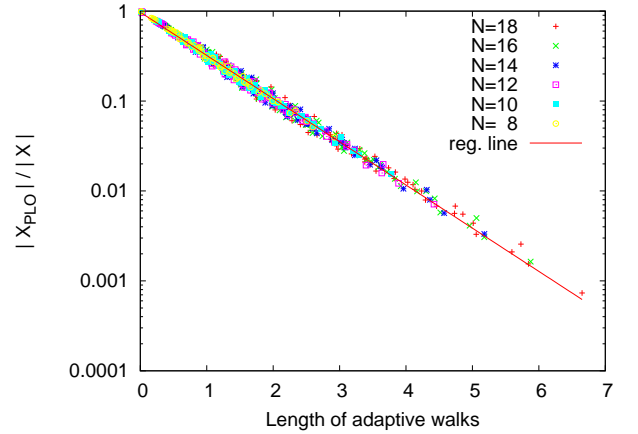


Figure 7: Scatter plot of the average density of PLO (to the size of the search space) vs. the average length of the Pareto adaptive walk for multiple N -values.

is $\log(y) = a \log(x) + b$ with $a = 1.059$ and $b = -6.536$. For such a log-log scale, the correlation is low. It is only possible to estimate the cardinality of the Pareto optimal set from the number of PLO with an order of magnitude. Nevertheless, there is a high positive correlation between the number of Pareto optimal solutions and the number of PLO.

5.3. Adaptive walk

In single-objective optimization, the length of adaptive walks, computed by means of a hill-climber, allows to estimate the average diameter of the local optima basins of attraction. That is, the number of local optima can be estimated when the whole search space cannot be enumerated exhaustively. Here, we define a multiobjective hill-climber, and we show that the

length of the corresponding adaptive walk is correlated with the number of PLO. However, a word of caution regarding the estimation by means of adaptive walks is the following: The height of the highest peak increases with the ruggedness (K), *i.e.* it decreases with the length of adaptive walks (Aguirre and Tanaka, 2004; Mathias et al., 2000; Skellett et al., 2005; Smith and Smith, 2000).

We define a very basic single solution-based *Pareto Hill-Climbing* (PHC). At each iteration of the PHC algorithm, the current solution is replaced by one dominating random neighboring solution. As a consequence, PHC stops on a PLO. The number of iterations, or steps, of the PHC algorithm is the length of the Pareto adaptive walk. For the ρMNK -landscapes investigated here, we consider the *1-bit-flip* neighborhood operator, which is directly related to a Hamming distance 1 over a bit string solution space.

We performed 10^3 independent PHC executions for each problem instance and $N \in \{8, 10, 12, 14, 16, 18\}$. The variation of the average length follows the opposite variation of the number of PLO. In order to show the link with the number of PLO, Figure 7 gives the scatter-plot of the average Pareto adaptive length *vs.* the logarithm of the average density of PLO for different instance sizes, as given in Table 2. The correlation is strong ($r = 0.997$), and the regression curve equation is: $|X_{PLO}|/|X| = C \cdot 2^{aL}$, where L is the average length of the adaptive walks, $C = 0.97$ and $a = -1.60$. For any ρMNK -landscapes parameter setting with a problem size $N \leq 18$, the average length of the Pareto adaptive walks can then give a precise estimation of the average number of PLO. When the adaptive length is short, the diameter of the basin of attraction associated with a PLO is short. This means that the distance between PLO decreases. Moreover, assuming that the volume of this basin is proportional to a power of its diameter, the number of PLO increases exponentially when the adaptive length decreases. This corroborates known results from single-objective NK-landscapes (Ochoa et al., 2008).

5.4. Properties *vs.* multi-modality for large-size problems

In this section, we study the number of PLO for large-size ρMNK -landscapes using the length of the adaptive walk as defined in the previous section. First, we analyze this number according to the problem dimension ($N \in \{18, 32, 64, 128\}$). Then, we give more insights on the problem difficulty, in terms of PLO, with respect to objective space dimension (M) and objective correlation (ρ).

We performed 10^3 independent PHC executions for each problem instance with $N \geq 18$. Figure 8 shows the average length of the Pareto adaptive walks for different landscapes according to the parameters given in Table 2. Whatever the objective space dimension and the objective correlation, the length of the adaptive walks increases linearly with the search space dimension N : $L = \alpha \cdot N$, with a slope $\alpha < 0.33$. The equation given in Section 5.3 allows to estimate the PLO density from the average length of the adaptive walk for small-size problems. Let us remind that $|X| = 2^N$. Following this estimation for large-size instances, the density of PLO decreases exponentially:

$|X_{PLO}|/|X| = C \cdot 2^{a\alpha N}$, such that $-0.6 < a\alpha < 0$. The coefficient $a\alpha$ is higher than -1 , so that the number of PLO increases exponentially: $|X_{PLO}| = C \cdot 2^{(1+a\alpha)N}$, with $0.4 < 1 + a\alpha < 1$. We can then reasonably conclude that the size of the Pareto optimal set grows exponentially as well (Section 5.2). The slope of the Pareto adaptive length increase is related to the objective space dimension (M) and to the objective correlation (ρ). The higher the number of objective functions, the smaller the slope. As well, the higher the objective correlation, the smaller the slope.

Figure 8 (bottom) gives a qualitative comparison for given problem sizes ($N = 64$ and $N = 128$). Let us consider an arbitrary adaptive walk length of 10. For ρMNK -landscapes with $N = 64$ and $K = 4$, this length corresponds approximately to the following parameter tuples at the same time: $(\rho = -0.4, M = 2)$, $(\rho = 0.3, M = 3)$, and $(\rho = 0.7, M = 5)$. As well, for $N = 128$ and $K = 4$, we have $(\rho = -0.9, M = 2)$, $(\rho = -0.1, M = 3)$, and $(\rho = 0.3, M = 5)$. Still assuming that the efficiency of a Pareto-based local search algorithm is closely related to the density of PLO, a problem with a small objective space dimension but a negative objective correlation can be more difficult to solve than another problem with many positively correlated objectives.

5.5. Discussion

This section gives one of the first analysis related to the local optima of a MoCO problem, based on the Pareto dominance relation. We first focused on small-size problems with a study of the number of PLO by complete enumeration. Like in single-objective optimization, the number of PLO increases with the degree of non-linearity of the problem (epistasis). However, the number of objective functions and the objective correlation have a stronger influence. Moreover, our results show that the cardinality of the Pareto optimal set clearly increases with the number of PLO. We proposed a Pareto adaptive walk, associated with a Pareto hill-climber, to estimate the number of PLO for small-size problems. Next, for large-size instances, the length of such Pareto adaptive walk can give a measure related to the difficulty of a MoCO problem. We show that this measure increases exponentially with the problem size. A problem with a small number of negatively correlated objectives gives the same degree of multi-modality, in terms of Pareto dominance, than another problem with a high objective space dimension and a positive objective correlation. This feature should be taken into account when dealing with large objective space dimension in *many-objective optimization* problems (Wagner et al., 2007).

6. Summary and future works

A qualitative summary of the results from the paper is given in Table 3. It provides a complementary view of the precise measures reported in the previous sections, and makes a stronger link with the following local search design issues.

- (i) **Archiving.** The approximation set manipulated by the algorithm (*i.e.* the archive), is directly related to the cardinality of the Pareto optimal set. The number of Pareto optimal solutions grows exponentially with the problem size,

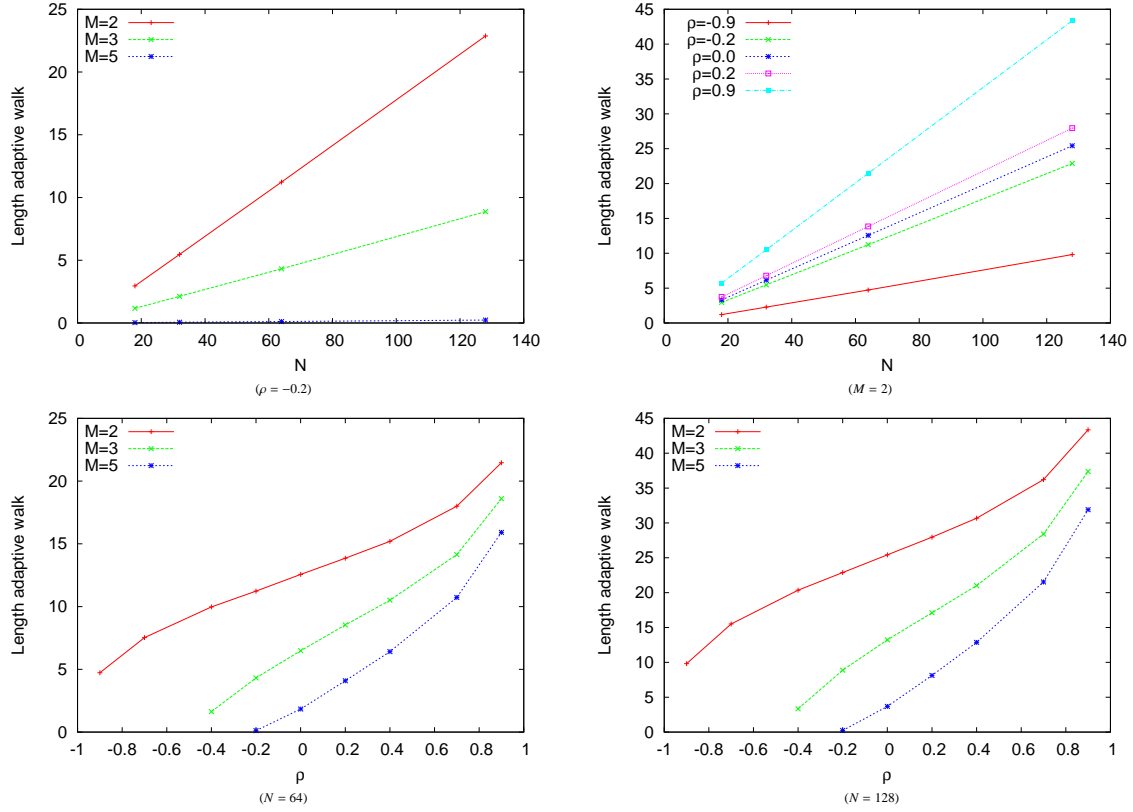


Figure 8: Average length of the Pareto adaptive walk according to problem size (N) for $K = 4$ and $\rho = -0.2$ (top-left) and for $K = 4$ and $M = 2$ (top-right). Average length of the Pareto adaptive walk according to objective correlation (ρ) for $K = 4$ and $N = 64$ (bottom-left) and for $K = 4$ and $N = 128$ (bottom-right).

the number of objectives and with the degree of conflict between the objectives. Moreover, it decreases with the non-linearity. As a consequence, the archive size should be limited as soon as the problem under consideration is large, and with many conflicting objectives. A review of existing bounded archiving techniques are provided by Knowles and Corne (2004), and more recently by López-Ibáñez et al. (2011).

(ii) **Scalar approach.** Scalar approaches are only able to find a subset of Pareto optimal solutions, known as supported solutions. The proportion of supported solutions in the Pareto optimal set decreases with the number of objectives, and their degree of conflict. Scalar approaches should then be well-performing when there is few correlated objectives. Of course, this always relies on the efficiency of the single-objective algorithm used in the scalar approach, and on the number of Pareto optimal solutions that are likely to be found, or approximated.

(iii) **Pareto-based approach.** The number of local optima, measured in terms of Pareto optimality, grows exponentially with the number of objectives, their degree of conflict, and also with the problem size. As opposed to its impact on the cardinality of the Pareto optimal set, the degree of non-linearity makes the number of Pareto local optima grows. Therefore, if we assume that the algorithm effi-

ciency is largely influenced by local optima, as in single-objective optimization, the performance of Pareto-based search approaches is higher when the problem is large, and aims at optimizing many conflicting objective functions. However, the performance quality of such approaches is with no doubt associated with an increase of the computational cost, particularly when a time-consuming archiving technique is used to bound the approximation set. This result supports previous conclusions from Paquete and Stützle (2006), where an experimental analysis focused on algorithmic components is conducted on the biobjective QAP.

Of course, other transversal views of Table 3, based *e.g.* on the problem properties, are also possible. The summary attempts to give a “big picture” of the multi-dimensional features that make local search algorithms efficient for the problem under study. All these results show that no expectation on the performance of multiobjective local search algorithms can be drawn without taking the problem properties into account very precisely. Indeed, it has now become clear that the number of objectives is one of the key issue to explain a problem complexity (Wagner et al., 2007), but we also pointed out that the objective correlation is at least as important. Multiobjective fitness landscape analysis plays a central role to explain the performance of local search algorithms, and to design more efficient methods, that suit better the problem features.

Table 3: Guidelines for the design of multiobjective local search algorithms.

Landscape features	Problem properties				Suggestion for the design of local search
	N	K	M	ρ	
Cardinality of the Pareto optimal set					+ limited-size archive - unbounded archive
Proportion of supported solutions					+ high efficiency of scalar approach - low efficiency of scalar approach
Number of Pareto local optima					+ low efficiency of Pareto approach - high efficiency of Pareto approach

A similar study, as the one led in the paper, would allow to better understand the structure of the landscape for other multiobjective combinatorial optimization problems. Moreover, additional work is required in order to compute, or rather estimate, the impact of the problem size on the proportion of supported solutions. Another perspective is to study the influence of the epistatic interactions, possibly defined with a different structure on every objective function, on the number of Pareto optimal solutions, and the proportion of supported solutions. Indeed, as studied originally by Aguirre and Tanaka (2007), a different epistasis structure *per* objective may change significantly the features of the landscape. At last, we know that the Pareto optimal set structure, as well as the number and the distribution of local optima have a strong impact on the performance of multiobjective search algorithms, but it is not yet clear how they exactly affect the search. The next step is to more precisely classify the relationship between problem structure and algorithm behavior. For sure, this open issue will constitute one of the main challenge in the emerging field of fitness landscape analysis for multiobjective combinatorial optimization.

Acknowledgments. The authors are grateful to the anonymous referees for their valuable and constructive comments that largely contributed to improve the quality of the paper.

References

- Aguirre, H. E., Tanaka, K., 2004. Insights on properties of multiobjective MNK-landscapes. In: IEEE Congress on Evolutionary Computation (CEC 2004). pp. 196–203.
- Aguirre, H. E., Tanaka, K., 2007. Working principles, behavior, and performance of MOEAs on MNK-landscapes. *European Journal of Operational Research* 181 (3), 1670–1690.
- Borges, P., Hansen, M., 1998. A basis for future successes in multiobjective combinatorial optimization. Tech. Rep. IMM-REP-1998-8, Institute of Mathematical Modelling, Technical University of Denmark, Lyngby, Denmark.
- Deb, K., 2001. *Multi-Objective Optimization using Evolutionary Algorithms*. John Wiley & Sons, Chichester, UK.
- Ehrgott, M., 2005. *Multicriteria optimization*, 2nd Edition. Springer.
- Ehrgott, M., Gandibleux, X., 2004. Approximative solution methods for multiobjective combinatorial optimization. *TOP* 12 (1), 1–89.
- Garrett, D., Dasgupta, D., 2007. Multiobjective landscape analysis and the generalized assignment problem. In: *Learning and Intelligent Optimization (LION 4)*. Springer-Verlag, Berlin, Heidelberg, pp. 110–124.
- Garrett, D., Dasgupta, D., 2009. Plateau connection structure and multiobjective metaheuristic performance. In: *Congress on Evolutionary Computation (CEC 2009)*. pp. 1281–1288.
- Hotelling, H., Pabst, M. R., 1936. Rank correlation and tests of significance involving no assumptions of normality. *Annals of Mathematical Statistics* 7, 29–43.
- Kauffman, S. A., 1993. *The Origins of Order*. Oxford University Press, New York, USA.
- Knowles, J., Corne, D., 2002. Towards landscape analyses to inform the design of a hybrid local search for the multiobjective quadratic assignment problem. In: *Soft Computing Systems: Design, Management and Applications*. Vol. 2002. IOS Press, Amsterdam, The Netherlands, pp. 271–279.
- Knowles, J., Corne, D., 2004. Bounded Pareto archiving: Theory and practice. In: *Metaheuristics for Multiobjective Optimisation*. Vol. 535 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Ch. 2, pp. 39–64.
- Liefvooghe, A., Humeau, J., Mesmoudi, S., Jourdan, L., Talbi, E.-G., 2012. On dominance-based multiobjective local search: design, implementation and experimental analysis on scheduling and traveling salesman problems. *Journal of Heuristics* 18, 317–352.
- López-Ibáñez, M., Knowles, J. D., Laumanns, M., 2011. On sequential online archiving of objective vectors. In: *6th International Conference on Evolutionary Multi-Criterion Optimization (EMO 2011)*. Vol. 6576 of *Lecture Notes in Computer Science*. Springer, pp. 46–60.
- Mathias, K. E., Eshelman, L. J., Schaffer, J. D., 2000. Niches in NK-landscapes. In: *Proceedings of the Sixth Workshop on Foundations of Genetic Algorithms (FOGA 2000)*. pp. 27–46.
- Merz, P., 2004. Advanced fitness landscape analysis and the performance of memetic algorithms. *Evolutionary Computation* 12 (3), 303–325.
- Mote, J., Olson, I. M. D. L., 1991. A parametric approach to solving bicriterion shortest path problems. *European Journal of Operational Research* 53 (1), 81–92.
- Ochoa, G., Tomassini, M., Verel, S., Darabos, C., 2008. A study of NK landscapes' basins and local optima networks. In: *Genetic and Evolutionary Computation Conference (GECCO 2008)*. ACM, pp. 555–562.
- Paquete, L., Schiavinotto, T., Stützle, T., 2007. On local optima in multiobjective combinatorial optimization problems. *Annals of Operations Research* 156 (1), 83–97.
- Paquete, L., Stützle, T., 2006. A study of stochastic local search algorithms for the biobjective QAP with correlated flow matrices. *European Journal of Operational Research* 169 (3), 943–959.
- Paquete, L., Stützle, T., 2007. Stochastic local search algorithms for multiobjective combinatorial optimization: A review. In: *Handbook of Approximation Algorithms and Metaheuristics*. Vol. 13 of *Computer & Information Science Series*. Chapman & Hall / CRC, pp. 1–29.
- Paquete, L., Stützle, T., 2009. Clusters of non-dominated solutions in multiobjective combinatorial optimization: An experimental analysis. In: *Multiobjective Programming and Goal Programming*. Vol. 618 of *Lecture Notes in Economics and Mathematical Systems*. Springer, pp. 69–77.
- Paquete, L., Stützle, T., 2010. On the performance of local search for the biobjective traveling salesman problem. In: *Advances in Multi-Objective Nature Inspired Computing*. *Studies in Computational Intelligence*. Springer, Berlin, Germany, Ch. 7, pp. 143–165.
- Serafini, P., 1987. Some considerations about computational complexity for multiobjective combinatorial problems. In: *Recent advances and historical development of vector optimization*. Vol. 294 of *Lecture Notes in Economics and Mathematical Systems*. Springer, pp. 222–232.
- Skellott, B., Cairns, B., Geard, N., Tonkes, B., Wiles, J., 2005. Maximally

- rugged NK landscapes contain the highest peaks. In: Genetic and Evolutionary Computation Conference (GECCO 2005). pp. 579–584.
- Smith, R. E., Smith, J., 2000. New methods for tunable, random landscapes. In: Proceedings of the Sixth Workshop on Foundations of Genetic Algorithms (FOGA 2000). pp. 47–67.
- Verel, S., Liefvooghe, A., Jourdan, L., Dhaenens, C., 2011a. Analyzing the effect of objective correlation on the efficient set of MNK-landscapes. In: Learning and Intelligent Optimization (LION 5). Vol. 6683 of Lecture Notes in Computer Science. Springer, Rome, Italy, pp. 238–252.
- Verel, S., Liefvooghe, A., Jourdan, L., Dhaenens, C., 2011b. Pareto local optima of multiobjective NK-landscapes with correlated objectives. In: 11th European Conference on Evolutionary Computation in Combinatorial Optimisation (EvoCOP 2011). Vol. 6622 of Lecture Notes in Computer Science. Springer, Torino, Italy, pp. 226–237.
- Wagner, T., Beume, N., Naujoks, B., 2007. Pareto-, aggregation-, and indicator-based methods in many-objective optimization. In: Fourth International Conference on Evolutionary Multi-Criterion Optimization (EMO 2007). Vol. 4403 of Lecture Notes in Computer Science. Springer-Verlag, Matsushima, Japan, pp. 742–756.
- Weinberger, E. D., 1996. NP completeness of kauffman’s N-k model, a tuneable rugged fitness landscape. Tech. Rep. 96-02-003, Santa Fe Institute.