

EXPLORING THE PARAMETER SPACE OF A SEARCH ALGORITHM

Katerina Tashkova

NELA Development Center, Železniki, Slovenia

katerina.taskova@ijs.si

Jurij Šilc, Peter Korošec

Computer Systems Department, Jožef Stefan Institute, Ljubljana, Slovenia

{jurij.silc; peter.korosec}@ijs.si

Abstract This paper addresses the problem of tuning the performance of a metaheuristic search algorithm with respect to its parameters. The operational behavior of typical metaheuristic search algorithms is determined by a set of control parameters, which have to be fine-tuned in order to obtain a best performance for a given problem. The principle challenge here is how to provide meaningful settings for an algorithm, obtained as result of better insight in its behavior. In this context, we discuss the idea of learning a model of an algorithm behavior by data mining analysis of parameter tuning results. The study was conducted using the Differential Ant-Stigmergy Algorithm as an example metaheuristic search algorithm.

Keywords: Data mining, Differential ant-stigmergy algorithm, Low-discrepancy sequences, Metaheuristic optimization, Parameter tuning.

1. Introduction

The research interest for metaheuristic search algorithms has been significantly growing in the last 20 years as a result of their efficiency and effectiveness to solve large and complex problems across different domains [2]. The state-of-the-art nature-inspired metaheuristic algorithms for high-dimensional continuous optimization include also algorithms inspired from the collective behavior of social organisms [13].

One such algorithm, which we will address in this paper, is the Differential Ant-Stigmergy Algorithm (DASA) initially proposed by Korošec [6], and further improved in [7]. DASA is inspired by the efficient self-

organizing behavior of ant colonies emerging from a pheromone-mediated communication, known as stigmergy [3]. Naturally, DASA can be classified within the Ant Colony Optimization (ACO) framework. However, the use of a continuous pheromone model in the form of Cauchy probability distribution with representation of the search space in the form of the so-called differential graph makes DASA dissimilar from the original ACO paradigm. The rationale behind DASA is in memorizing (via the pheromone model updates) the “move” in the search space that improves the current best solution and using it in further search.

As it is the case with most of the metaheuristic algorithms, the operational behavior of DASA is determined by a set of control parameters. In practice, these parameters have to be fine-tuned in order to obtain best performance of the algorithm. It can be quite inconvenient for the users as: (i) they usually do not have insight into the behavior of the algorithm, and (ii) even if a default setting exists, it may not be adequate for a specific instance or type of a problem. Moreover, parameter tuning is computationally expensive task.

The principle challenge here is how to provide meaningful (default) settings for DASA, obtained as result of better insight into the algorithm’s behavior. Furthermore, can we find optimal regions in DASA parameter space by analyzing the patterns in the algorithm’s behavior with respect to the problem characteristics? Related to this, we discuss the preliminary findings based on data mining analysis of parameter tuning results. More precisely, the parameter tuning task is approached by two-step procedure that combines a kind of experimental design with data mining analysis.

We use Sobol’ sequences [9] for even sampling of the algorithm parameter space to generate a large and diverse set of parameter settings. These are used as input to DASA to be tuned on a given function optimization problem. The performance of DASA on the given function optimization problem, in terms of function error, is captured at different time points for all sampled parameter settings. The data collected in the first step, DASA performance with corresponding parameter settings, is subject for intelligent data analysis, i.e., multi-target regression with Predictive Clustering Trees [1].

Parameter sampling combined with regression has been already used by Stoean et al. [10] for tuning metaheuristics: Latin hypercubes parameter sampling is combined with single-target regression with Support Vector Machines. Our approach modifies the former by replacing the Latin hypercube sampling by Sobol’ sequences, as the former is best suited in the case when a single parameter dominates the algorithm’s performance, while it should be used with care if there are interactions

among the sampled parameters [8]. Moreover, we define the regression task as multi-target regression, taking into account more than one target (in this case the function error at few time points) with the goal to find parameter settings for the given algorithm that will not only solve the problem but will also solve the optimization problem fastest.

The remainder of this paper is structured as follows. Section 2 introduces the differential ant-stigmergy algorithm. Then, Section 3 addresses the parameter tuning task and Section 4 presents the experimental evaluation with the results. After that, Section 5 discusses the idea of post-hoc analysis of parameter tuning by data mining. Finally, Section 6 summarizes this study and outlines possible directions for further work.

2. The Differential Ant-Stigmergy Algorithm

The version of DASA used in our experimental evaluation is described in details by Korošec et al. [7]. In principle, DASA relies on two distinctive characteristics, differential graph and continuous pheromone model. Here, we will briefly discuss these two characteristics and outline the main loop of the DASA search process.

First, DASA transforms the D -dimensional optimization problem into a graph-search problem. The corresponding differential graph is a directed acyclic graph obtained by fine discretization of the continuous parameters' offsets. The graph has D layers with vertices, where each layer corresponds to a single parameter. Each vertex corresponds to a parameter offset value that defines a change from the current parameter value to the parameter value in the next search iteration. Furthermore, each vertex in a given layer is connected with all vertices in the next layer. The set of possible vertices for each parameter depends on the parameter's range, the *discretization base*, b , and the maximal computer arithmetics precision, which defines the minimal possible offset value. Ants use these parameters' offsets to navigate through the search space. At each search iteration, a single ant positioned at a certain layer moves to a specific vertex in the next layer, according to the amount of pheromone deposited in the graph vertices belonging to this layer.

Second, DASA performs pheromone-mediated search that involves best-solution-dependent pheromone distribution. The amount of pheromone is distributed over the vertices according to the Cauchy Probability Density Function (CPDF) [8]. DASA maintains a separate CPDF for each parameter. Initially, all CPDFs are identically defined by a location offset set to zero and a scaling factor set to one. As the search process progresses, the shape of the CPDFs changes: CPDFs shrink and

stretch as the scaling factor decreases and increases, respectively, while the location offsets move towards the offsets associated with the better solutions. The search strategy is guided by three user-defined real positive factors: the *global scale increase factor*, s_+ , the *global scale decrease factor*, s_- , and the *pheromone evaporation factor*, ρ . In general, these three factors define the balance between exploration and exploitation in the search space. They are used to calculate the values of the scaling factor and consequently influence the dispersion of the pheromone and the moves of the ants.

Finally, the main loop of DASA consists of an iterative improvement of a temporary-best solution, performed by searching appropriate paths in the differential graph. The search is carried out by m ants, all of which move simultaneously from a starting vertex to the ending vertex at the last level, resulting in m constructed paths. Based on the found paths, DASA generates and evaluates m new candidate solutions. The best among the m evaluated solutions is preserved and compared to the temporary-best solution. If it is better than the temporary-best solution, the latter is replaced, while the pheromone amount is redistributed along the path corresponding to the path of the preserved solution and the scale factor is accordingly modified to improve the convergence. If there is no improvement over the temporary-best solution, then the pheromone distributions stay centered along the path corresponding to the temporary-best solution, while their shape shrinks in order to enhance the exploitation of the search space. If for some fixed number of tries all the ants only find paths composed of zero-valued offsets, the search process is restarted by randomly selecting a new temporary-best solution and re-initializing the pheromone distributions.

3. Parameter Tuning

To obtain the best possible performance on a given problem, one should consider a task specific tuning of the parameter setting for the optimization algorithm used. Determining the optimal parameters is an optimization task in itself, which is extremely computationally expensive. There are two common approaches for choosing parameters values: parameter tuning and parameter control. The first approach selects the parameter settings before running the optimization algorithm (and they remain fixed while performing the optimization). The second approach optimizes the algorithm's parameters along with the problem's parameters. Here, we will focus on the first approach, parameter tuning.

A detailed discussion and survey of parameter tuning methods is given by Eiben and Smit [4]. According to this survey, one way to approach

parameter tuning is by *sampling* methods. Sampling methods reduce the search effort by decreasing the number of investigated parameter settings as compared to the full factorial design: the basic full factorial design investigates 2^k parameter settings, subject to k parameters, each of which have 2 possible values; in the more general case, parameters can have arbitrary number of values; moreover, an increase in the number of investigated parameters means an exponential increase in the number of parameter settings to be tested. Two widely used sampling methods are Latin-squares [8] and Taguchi orthogonal arrays [11]. However, these are not the most robust sampling techniques, e.g., Latin-squares or Latin hypercube sampling is good in the case where one of the parameters dominates the algorithm's performance, while it should be used with care if there are interactions among the parameters.

Ultimately, we would like to find a sampling schema that will be able to detect the interactions among the parameters, will be independent from user-specified information regarding the particular parameter values to be considered (typical for factorial design), and will deliver small but representative sample of the parameter search space. The first two requirements are satisfied by the pure random sampling, but the last is not, as random sampling does not guarantee that the sampled values are evenly spread across the entire domain. The so-called low-discrepancy sequences were specially designed to fulfill all three requirements. Therefore, Sobol' sequences, a representative variation of low-discrepancy sequences introduced by Sobol' [9], was considered for sampling the parameter space of DASA in this study.

Sobol' sequences, sampled from a D -dimensional unit search space, are quasi-random sequences of D -tuples that are more uniformly distributed than uncorrelated random sequences of D -tuples. These sequences are neither random nor pseudo-random: they are cleverly generated not to be serially uncorrelated by taking into account which tuples in the search space have already been sampled. For a detailed explanation and overview of the schemas for generating Sobol' sequences, we refer to Press et al. [8]. The particular implementation of Sobol' sampling used in our analysis is based on the Gray code order [5].

4. Experimental Evaluation

Since data mining methods can only discover patterns actually present in the data, the dataset subject to analysis must be large enough and informative enough to contain these patterns, i.e., to describe different types of algorithm's behavior. For this reason, we decided to use a simple

test function, which matched this requirement and was used for building an example case model.

Therefore, the performance of DASA was evaluated on the Sphere function:

$$f(\vec{x}) = \|\vec{z}\|^2 + f(\vec{x}_{\text{opt}}),$$

where $\vec{z} = \vec{x} - \vec{x}_{\text{opt}}$ and \vec{x}_{opt} is optimal solution vector, such that $f(\vec{x}_{\text{opt}})$ is minimal. Function $f(x)$ is defined over D -dimensional real-valued search space x and is scalable with the dimension D . It has no specific value of its optimal solution (it is randomly shifted in x -space) and has an artificially chosen optimal value (it is randomly shifted in f -space). In this study, we considered the Sphere function with respect to two dimensions, $D = 20$ and $D = 40$.

The performance of DASA is dependent on the values of five parameters: three real-valued parameters that directly influence the search heuristic (s_+ , s_- , and ρ) and two integer-valued parameters (m and b). Therefore, we considered all of them for tuning DASA performance on the Sphere function for both search space dimensions, $D = 20$ and $D = 40$. Using the Gray-code-based Sobol' generator we generated 5000 parameter settings (5-tuples). Note that the Sobol' sampling generates numbers on the unit interval: in order to obtain the true parameter settings, we mapped these values on the predefined search range of parameter values. The latter for each of the five tuned parameters was defined as follows: $4 \leq m \leq 200$, $0 \leq \rho \leq 1$, $0 \leq s_+ \leq 1$, $0 \leq s_- \leq \rho$, and $2 \leq b \leq 100$. Moreover, the mapped values for the integer-valued parameters m and b were rounded to the closest integer value. Finally, due to implementation reasons, the upper bound of the global scale decrease factor s_- was actually limited by the value of the evaporation factor ρ .

In the next step, the performance of the Sobol' sampled parameter settings were tested on the Sphere benchmark function. Due to the stochastic nature of DASA, every parameter setting was used in a multiple-run experimental evaluation. Each run included $25,000 \times D$ function evaluations (FEs). The number of runs was set to 15.

The results gathered by the parameter tuning process are most often subjected to ordinal data analysis, which includes ranking of the different sampled parameter sets according to some calculated statistics, e.g., best or mean performance of the algorithm in some predefined number of runs [12]. In this case, performance of the algorithm is expressed in terms of the function error, i.e., the difference between the obtained and optimal function value. In order to find a setting that will be satisfactory in terms of convergence speed, we captured the error values at four different time

points, corresponding to $25 \times D$, $250 \times D$, $2,500 \times D$, and $25,000 \times D$ FEs.

The optimal performing parameter setting was chosen based on the median best performance over all runs aggregated over all time points for a given dimension ($D = 20$ and $D = 40$). A common approach is to use the mean performance, but we took the median in order to avoid the problems that the mean has when observing large variance in the function values across the runs. More precisely, given a function, an individual rank is assigned to every setting (out of 5000) for the four time points. A single final rank is calculated by ranking the sum of the four individual rankings assigned to the parameter settings. The best-ranked parameter setting for a given dimension defines instance of DASA referred to as DASA*.

The results of DASA tuning subject to ordinal data analysis are presented in Tables 1 and 2. Table 1 reports the tuned parameter settings for both DASA* instances. In addition, the default parameter setting for DASA from [7] is given as a reference for comparison. The corresponding instance is referred to as DASA°.

Table 1. Parameter settings for DASA° and DASA*.

Algorithm	Dimension	m	ρ	s_+	s_-	b
DASA°	20, 40	10	0.2	0.02	0.01	10
DASA*	20	5	0.324	0.201	0.289	6
	40	7	0.388	0.136	0.344	8

Results in Tables 2 represent the median values of the function errors, at the four time points, obtained by DASA* and DASA° instances for both dimensions. Note, that error value below 10^{-8} was treated as zero. The tables clearly show that DASA* is better than DASA°.

Table 2. Median values of the function errors for the Sphere function.

Algorithm	DASA°		DASA*	
	$D = 20$	$D = 40$	$D = 20$	$D = 40$
FEs				
$25 \times D$	16.7	18.3	2.53	9.31
$250 \times D$	0.0003	0.0021	0	0
$2,500 \times D$	0	0	0	0
$25,000 \times D$	0	0	0	0

5. Data-Mining Analysis

Parameter tuning of an algorithm leads to a better performance, however it is a computationally expensive and problem-dependent task. Considering this, the idea is to extend the simple tuning that delivers a single parameter set and analyze the gathered data in an intelligent way. The intelligent analysis can extract patterns (regularities) in the explored parameter space that define a specific behavior of DASA. To this end, data mining methods for automated discovery of patterns in data can be used. As data mining methods can only discover patterns that are present in the data, the dataset subject to analysis must be large enough and informative enough to contain these patterns, i.e., to describe different types of algorithm's behavior. Related to this, we considered a data mining approach on a representative example, i.e., error model of the Sphere function.

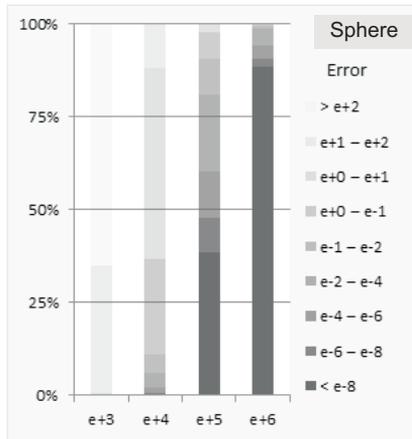


Figure 1. Median error distributions for the Sphere function in the case of $D = 40$.

To begin with, consider the graph in Fig. 1 that visualizes the Sphere function error distribution. The graph depicts the distribution of the median error values obtained by 5000 parameter settings at four different points for $D = 40$. As we are more concerned with the practical significance between a large and a small error value than the statistical significant difference between two actual error values, the error values are discretized into nine intervals, each of which is represented with a color chosen according to the error magnitude between black (error below 10^{-8}) and white (error above 10^2). The graph clearly shows that the sampled settings determine different DASA performance. As evident, there is a big cluster of parameter settings that solve this function

to the aimed accuracy (error below 10^{-8}) in the given time budget (10^6 FEs). Moreover, subset of this cluster solves the function for an order of magnitude less FEs. Our aim, therefore, is to find a (common) description of this cluster, in terms of DASA parameter relations, that represents a good behavior of DASA (as well as what parameter relations lead to a bad DASA performance).

For this purpose, we formulated the problem as a predictive modeling task using decision trees to model the function error values in terms of the values of DASA parameters. Since the function error variables are continuous, the task at hand is a regression task. Furthermore, as our goal is to model the behavior of DASA at all time points simultaneously, the problem at hand is then a multi-target regression. To this end, we used Predictive Clustering Trees (PCTs) which are implemented in Clus system [1]. In this case, the median error at the four time points define the four target attributes considered for modeling (with the PCT) in dependence from the descriptive attributes, i.e., the function dimension and the five DASA parameters. The resulting dataset is composed of 10,002 rows described with the 5001 parameter settings (including the default setting) of DASA applied to the two different dimensions, and 10 columns corresponding to the 10 attributes.

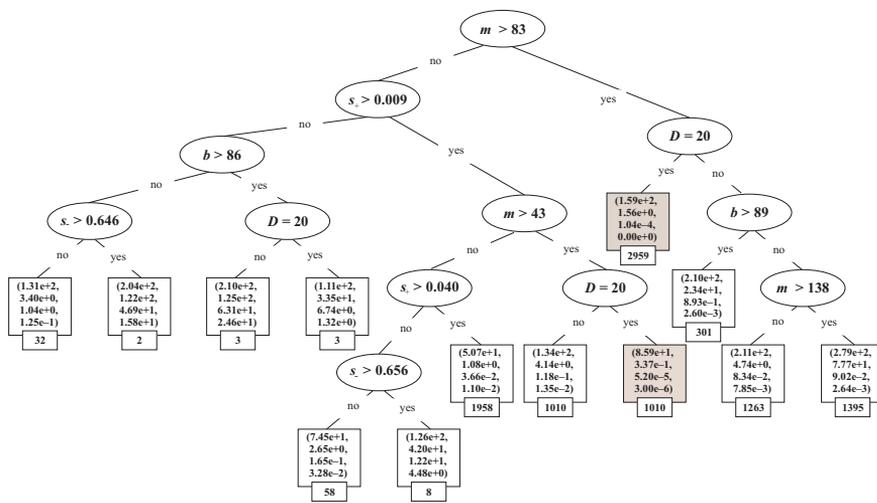


Figure 2. Predictive clustering tree representation of the error model for the Sphere function.

Figure 2 presents a PCT model for the Sphere function error. Each internal node of the tree contains a test on a descriptive attribute, while the leaves store the model predictions for the function error, i.e., a tu-

ple of four values. The predictions are calculated as the mean values of the corresponding error values for the data instances belonging to the particular leaf (represented by a box). In fact, each leaf identifies one cluster of data instances (the size of the cluster is the value in the small box). The predictive performance of the model was assessed with 10-cross-validation. Note that this particular model was learned on the complete dataset subject to constraints on the maximal tree size of 25 nodes. We did this because the original model contained 1643 nodes (of which 822 leaves) and despite its better predictive performance, both training and testing, it was not comprehensible; aiming for an explanatory model, small and comprehensible, we considered the smaller tree obtained with the limitation of the size. The predictive performance of both models in terms of Root Relative Mean Squared Error (RRMSE) and Pearson Correlation Coefficient (PCC) are given in Table 3. Note that RRMSE represents the relative error with respect to the mean predictor performance, while PCC represents the linear correlation between the data and the model predictions. Good models have RRMSE values closer to 0 and PCC closer to 1.

Table 3. Model performance with respect to RRMSE and PCC.

Performance	PCT nodes	Measure	FEs				Mean
			$2.5e+2 \times D$	$2.5e+3 \times D$	$2.5e+4 \times D$	$2.5e+5 \times D$	
Training	1643	RRMSE	0.166	0.373	0.487	0.472	0.396
		PCC	0.986	0.928	0.873	0.882	0.843
	25	RRMSE	0.408	0.679	0.562	0.546	0.557
		PCC	0.913	0.734	0.827	0.837	0.689
Testing	1643	RRMSE	0.219	0.638	1.019	1.053	0.806
		PCC	0.976	0.777	0.304	0.234	0.426
	25	RRMSE	0.450	0.817	1.039	1.055	0.875
		PCC	0.893	0.584	0.246	0.218	0.312

The model in Fig. 2 outlines 13 clusters of data instances, of which two (depicted with light-gray boxes) represent a good DASA performance. According to this model, the number of ants, m , is the most important DASA parameter for its performance on the Sphere function. More precisely, if $m > 83$, independent of the values of the other parameters, DASA solves the 20-dimensional functional problem for the given time budget. Moreover, if $m \leq 83$ another DASA parameters become important as well. For example, if $43 < m \leq 83$ and $s_+ > 0.009$ and $D = 20$ then DASA solves the function with error $3 \cdot 10^{-6}$, while the pattern $m \leq 43$ and $s_+ \leq 0.040$ and $s_- > 0.656$ describes a poor DASA performance regardless of the function dimension. An interesting fact is

that, the evaporation factor is not essential for DASA performance on the Sphere function. Moreover, the model also shows that is more difficult to describe the behavior of DASA for the 40-dimensional function problem than the 20-dimensional one.

Finally, note that the training performance (learned on the complete dataset) of the model in terms of the error and the correlation coefficient is best for the first target, while it gets worse with respect to the other three targets (see Table 3). This is especially significant if we take into account the testing performance of the model estimated with 10-cross-validation. However, the training performance is acceptable in our case, as we are interested in understanding the behavior of DASA and not aiming to obtain a model for prediction.

6. Conclusions

The principle challenge of metaheuristic design is providing a default algorithm configuration, in terms of parameter setting, that will perform reasonably well in general (problem) case. However, while it is a good initial choice, the default algorithm configuration may result in low quality solutions on a specific optimization problem. In practise, the algorithms parameters have to be fine-tuned in order to obtain best algorithm's performance for the problem at hand, leading to the computational expensive task of parameter tuning. So, if the tuning task is unavoidable, the question is: can we use the results from the parameter tuning to extract some knowledge about the algorithm's behavior?

Related to this, the study focused on the problem of tuning the performance of the Differential Ant-Stigmergy Algorithm (DASA) with respect to its parameters. As it is the case with most of the metaheuristic algorithms, the operational behavior of DASA is determined by a set (five) of control parameters. The existing default setting of DASA parameters [7] is obtained by experimentation with both real and benchmark optimization problems, but not as a result of some systematic evaluation. Furthermore, there is no deeper understanding of the impact of a particular parameter or parameters relations on the performance of DASA. In this context, we performed an systematic evaluation of DASA performance obtained by solving the Sphere function optimization problem with 5000 Sobol' sampled DASA parameter settings regarding two dimensions, 20 and 40. Furthermore, we discussed the idea of learning a model of DASA behavior by data mining analysis of the parameter tuning results. In this context, we formulated the problem as multi-target regression and applied predictive clustering trees for learning a model of DASA behavior with respect to the function error performance. The ob-

tained model revealed that the parameter denoting number of ants is the most important parameter for DASA performance on the 20-dimensional function problem. On the other hand, the evaporation factor is not essential for DASA performance on the Sphere function.

Further work will focus on additional experimental evaluation and data mining analysis of data with respect to more complex functions problems. This idea can be further extended to building models of DASA behavior that will include the optimization problem characteristics (such as multimodality, separability, and ill-conditioning) as descriptive attributes as well. The latter can provide insights on how to configure DASA performance with respect to the type of the optimization problem. Moreover, these insights can serve as a valuable information for improvement of DASA design.

References

- [1] H. Blockeel and J. Struyf. Efficient algorithms for decision tree cross-validation. *J. Mach. Learn. Res.*, 3:621–650 (2002).
- [2] C. Blum and A. Roli. Metaheuristics in combinatorial optimization: overview and conceptual comparison. *ACM Comput. Sur.*, 35(3):268–308 (2003).
- [3] E. Bonabeau, M. Dorigo and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 1999.
- [4] A. E. Eiben and S. K. Smit. Parameter tuning for configuring and analyzing evolutionary algorithms. *Swarm Evol. Comput.*, 1(1):19–31 (2011).
- [5] S. Joe and F. Y. Kuo. Constructing Sobol sequences with better two-dimensional projections. *SIAM J. Sci. Comput.*, 30(5):2635–2654 (2008).
- [6] P. Korošec. Stigmergy as an approach to metaheuristic optimization. Doctoral Dissertation, Jožef Stefan International Postgraduate School, Ljubljana, 2006.
- [7] P. Korošec, J. Šilc and B. Filipič. The differential ant-stigmergy algorithm. *Inform. Sciences*, 192(1):82–97 (2012).
- [8] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery. *Numerical Recipes, 2nd edition*. Cambridge University Press, 1992.
- [9] I. M. Sobol'. Distribution of points in a cube and approximate evaluation of integrals. *USSR Comput. Maths. Math. Phys.*, 7(4):86–112 (1967).
- [10] R. Stoean, T. Bartz-Beielstein, M. Preuss, and C. Stoean. A Support Vector Machine-Inspired Evolutionary Approach for Parameter Setting in Metaheuristics Research Center CIOP (Computational Intelligence, Optimization and Data Mining). CIOP Technical Report. Cologne University of Applied Science, Faculty of Computer Science and Engineering Science, 2009.
- [11] G. Taguchi and T. Yokoyama. *Taguchi Methods: Design of Experiments*. ASI Press, 1993.
- [12] E-G. Talbi. *Metaheuristics: From Design to Implementation*. John Wiley and Sons, 2009.
- [13] X.-S. Yang. *Nature-Inspired Metaheuristic Algorithms*. Luniver Press, 2008.