

Chapter 1

Metaheuristics: Intelligent Problem Solving

Marco Caserta and Stefan Voß

Abstract Metaheuristics support managers in decision making with robust tools providing high quality solutions to important problems in business, engineering, economics and science in reasonable time horizons. While finding exact solutions in these applications still poses a real challenge despite the impact of recent advances in computer technology and the great interactions between computer science, management science/operations research and mathematics, (meta-) heuristics still seem to be the methods of choice in many (not to say most) applications. In this chapter we give some insight into the state of the art of metaheuristics. It focuses on the significant progress regarding the methods themselves as well as the advances regarding their interplay and hybridization with exact methods.

1.1 Introduction

The use of heuristics and metaheuristics to solve real world problems is widely accepted within the operations research community. It is well-known that the great majority of complex real world decision problems, when modeled as optimization problems, belong to the class of \mathcal{NP} -hard problems. This implies that exact approaches are doomed to fail when dealing with large scale instances, whether they arise from business, engineering, economics or science. Today, decision making processes are increasingly complex and more encompassing, in the sense that more decision variables are used to model complex situations and more input data and parameters are available to capture the complexity of the problems themselves.

Marco Caserta · Stefan Voß
Institute of Information Systems (Wirtschaftsinformatik), University of Hamburg,
Von-Melle-Park 5, 20146 Hamburg, Germany
e-mail: {marco.caserta, stefan.voss}@uni-hamburg.de

The inherent complexity of real world optimization problems, though, should be interpreted in the light of what complexity analysis really means: on the one hand, the fact that a problem belongs to the class of \mathcal{NP} -hard problems implies that there is no knowledge of an algorithm capable of solving the problem itself to optimality in polynomial time with respect to its input size (and many believe that there will never be); on the other hand, it is worth remembering that complexity analysis provides a worst case scenario, i.e., it indicates that, in the worst case, with the growth of the input size, the algorithm will require more and more time/steps to provide a definite answer. However, it is also worth noting that practitioners have observed that it is possible to design *ad-hoc* algorithms which, while not guaranteeing optimal or near-optimal solutions for the whole set of possible instances of a given problem (and, in many cases, not even a simple, single answer can be guaranteed at all), they do provide near-optimal solutions “most” of the times. This is exactly the goal of a metaheuristic designer (analyst), namely, to design an algorithm for which, even though no guarantee about the worst case scenario can be offered, a certain degree of confidence about the performance of the algorithm “most of the times” can still be asserted.

Consequently, the real challenge of the metaheuristic expert is not only to objectively measure the algorithm in terms of solution quality and computational time over problems for which no optimal solution is known (which, in itself, can be a challenging task due to the lack of any benchmark) but also to use sound quantitative methods and techniques to assert the robustness of the algorithm over a wide spectrum of instance types, hence enhancing the usability of the algorithm itself by industry decision makers, e.g., as “optimization modules” within decision support systems.

A first trade-off emerges here: on the one hand, it is highly desirable to design “general purpose” metaheuristics, which do not require problem specific knowledge and can readily be applied to a wide spectrum of problem classes. This has been the line of research of the last decades, during which a number of general purpose metaheuristic paradigms have been proposed, e.g., simulated annealing, genetic algorithms, tabu search, ant colony, etc. The main argument in favor of such paradigms is exactly their general applicability upon a large set of problems, without requiring major re-design or any in-depth knowledge of the problem to be tackled. Consequently, general paradigms seem especially suited for practitioners, who are interested in getting a solution to the problem without investing a huge amount of time in understanding the mathematical properties of the model and in implementing tailor-made algorithms. However, most successful metaheuristics applied to specific problems are also tailored to the problem, or fine-tuned. On the other hand, it has been observed that general purpose metaheuristics are outperformed by hybrid algorithms, which usually are algorithms that combine mathematical programming techniques with metaheuristic-based ideas. Algorithms of this class are tailor-made and specifically designed to exploit the mathematical properties of the problem at hand. While such approaches are

able to provide enhanced performance, an obvious consequence is a reduction in the usability of the algorithm itself. In general, a tailor-made algorithm can be used only for a specific class of problems and, often, the underlying ideas cannot easily be extended towards operating on a different class of problems.

The discussion on metaheuristic based algorithms is further complicated by what has been observed in recent years: these algorithms in general seem to provide varying performance depending upon the sensitivity (skills, expertise, ingenuity, etc.) of the designer in algorithmic fine tuning. In order to maximize algorithmic performance, an instance specific fine tuning might be required. The variability in the results of a metaheuristic presents at least two major drawbacks: (i) On one hand, the issue of *reproducibility* of results arises: It is common knowledge that a proposed algorithm implemented by two different researchers can lead to altogether different results, depending upon factors such as implementation skills, special usage of data structures, and ability in parameter settings, among others. For this reason, it is hard to compare and thoroughly assess a metaheuristic and its performance. (ii) On the other hand, a problem of *performance maximization* is envisioned: One of the commonalities of virtually all proposed metaheuristic paradigms is that they are characterized by a considerable number of parameters, whose value(s) strongly affect the overall performance of the algorithm itself. It is common knowledge that the fine tuning of algorithmic parameters is not only problem-specific, but even instance-specific, i.e., even for a given problem, parameter values should be fine-tuned and adjusted according to instance specific information (e.g., instance size, distribution of its values, etc.).

The purpose of this paper is to provide a survey of the general field of metaheuristics. While we cannot be fully comprehensive in a single paper, in line with the above remarks, some of the issues addressed in this paper are:

- In light of the well-known *no-free-lunch-theorem* [128], which basically states that, on average, no algorithm outperforms all the others, one might wonder what strategy to pursue, i.e., whether the goal of a researcher in the field should be to develop a better general framework able to effectively solve a wider spectrum of problems or, conversely, to tackle each individual problem separately, by designing tailor-made algorithms that fully exploit the mathematical structure and properties of each problem. A recent line of research in the metaheuristic field is concerned with the design of *hybrid* algorithms, where the term hybrid can indicate either the combination of different metaheuristics or the intertwined usage of metaheuristic features with mathematical programming techniques. Consequently, a trade-off between *re-usability* and *performance* of an algorithm arises.
- As highlighted before, no heuristic can guarantee high quality solutions over all possible instances of a given problem class. However, it is at least desirable to present a *robust* behavior over a spectrum of instances belonging to the same problem class. One key factor that seems to have a strong impact on the algorithmic performance is the fine tuning of the algorithm itself. Since the behavior of a metaheuristic is affected by its parameters,

one might wonder how to select a good set of parameter values. An important topic in metaheuristic design is, therefore, the identification and development of techniques for the *fine tuning of algorithmic parameters*.

- Due to the stochastic behavior of some metaheuristics and the lack of commonly accepted and adopted techniques for the evaluation of algorithmic performance, given two algorithms designed to tackle the same class of problems, it is not always possible to “rank” such algorithms in terms of their performance. One direct consequence is that there is still no clear understanding about which features are really successful and under which circumstances. In other words, unless clear standards about metaheuristics are defined, it will be really hard to have a full grasp of, first, which algorithms are better than others and, second, what is the real contribution of each feature of the algorithm upon the overall performance.

The structure of this paper is as follows: We first present general concepts about heuristics and metaheuristics, seen from an operations research perspective. Next, we illustrate some findings from recent research about hybridization, seen as a combination of exact approaches and mathematical programming techniques with metaheuristic frameworks. We also include here some recent contributions about metaheuristic design, i.e., a collection of ideas and thoughts about what should influence which features of a metaheuristic paradigm to be included in the algorithm (e.g., fitness landscape evaluation). Next, we focus on the important issue of metaheuristics fine tuning and calibrations, by introducing some quantitative methods drawn from statistics that have been employed with this goal in mind. We also present some thoughts on the ongoing debate on metaheuristic assessment, i.e., how an objective evaluation of the performance of a metaheuristic can be carried out in order to increase objectivity of the evaluation and reproducibility of the results. Next, we mention some optimization software libraries especially focused on the implementation of general heuristic and metaheuristic frameworks. The development of software libraries for metaheuristics is perceived as a key factor in the emerging of standards in the field. Finally, the last section presents some concluding remarks.

Earlier survey papers on metaheuristics include [19, 121, 122].¹ The general concepts have not become obsolete, and many changes are mainly based upon an update to most recent references. A handbook on metaheuristics is available describing a great variety of concepts by various authors in a comprehensive manner [59].

¹ Here we occasionally rely on [121] and [122] without explicitly quoting at appropriate places for not “disturbing” the readability.

1.2 Basic Concepts and Discussion

The basic concept of heuristic search as an aid to problem solving was first introduced by [93]. A *heuristic* is a technique (consisting of a rule or a set of rules) which seeks (and hopefully finds) *good* solutions at a reasonable computational cost. A heuristic is *approximate* in the sense that it provides (hopefully) a good solution for relatively little effort, but it does not guarantee optimality.

Heuristics provide simple means of indicating which among several alternatives seems to be best. That is, “heuristics are criteria, methods, or principles for deciding which among several alternative courses of action promises to be the most effective in order to achieve some goal. They represent compromises between two requirements: the need to make such criteria simple and, at the same time, the desire to see them discriminate correctly between good and bad choices. A heuristic may be a *rule of thumb* that is used to guide one’s action.” [91]

Greedy heuristics are simple iterative approaches available for any kind of (e.g., combinatorial) optimization problem. A good characterization is their *myopic* behavior. A greedy heuristic starts with a given feasible or infeasible solution. In each iteration there is a number of alternative choices (*moves*) that can be made to transform the solution. From these alternatives which consist in fixing (or changing) one or more variables, a *greedy choice* is made, i.e., the best alternative according to a given measure is chosen until no such transformations are possible any longer.

Usually, a greedy *construction heuristic* starts with an incomplete solution and completes it in a stepwise fashion. Savings and dual algorithms follow the same iterative scheme: Dual heuristics change an infeasible low cost solution until reaching feasibility, savings algorithms start with a high cost solution and realize the highest savings as long as possible. Moreover, in all three cases, once an element is chosen this decision is (usually) not reversed throughout the algorithm, it is kept.

As each alternative has to be measured, in general we may define some sort of *heuristic measure* (providing, e.g., some priority values or some ranking information) which is iteratively followed until a complete solution is build. Usually this heuristic measure is applied in a greedy fashion.

For heuristics we usually have the distinction between finding initial feasible solutions and improving them. In that sense we first discuss local search before characterizing metaheuristics.

1.2.1 Local Search

The basic principle of local search is to successively alter solutions *locally*. Related transformations are defined by neighborhoods which for a given solution

include all solutions that can be reached by one move. That is, neighborhood search usually is assumed to proceed by moving iteratively from one solution to another one by performing some sort of operation. More formally, each solution of a problem has an associated set of neighbors called its *neighborhood*, i.e., solutions that can be obtained by a single operation called transformation or move. Most common ideas for transformations are, e.g., to add or drop some problem specific individual components. Other options are to exchange two components simultaneously, or to swap them. Furthermore, components may be shifted from a certain position into other positions. All components involved within a specific move are called its elements or attributes.

Moves must be evaluated by some *heuristic measure* to guide the search. Often one uses the implied change of the objective function value, which may provide reasonable information about the (local) advantage of moves. Following a greedy strategy, *steepest descent* (SD) corresponds to selecting and performing in each iteration the best move until the search stops at a local optimum. Obviously, savings algorithms correspond to SD.

As the solution quality of local optima may be unsatisfactory, we need mechanisms that guide the search to overcome local optimality. For example, a metaheuristic strategy called iterated local search is used to iterate/restart the local search process after a local optimum has been obtained, which requires some perturbation scheme to generate a new initial solution (e.g., performing some random moves). Of course, more structured ways to overcome local optimality may be advantageous.

A general survey on local search can be found in [1] and the references from [2]. A simple template is provided by [116].

Despite the first articles on this topic already being in the 1970s (cf. Lin and Kernighan [82]), a variable way of handling neighborhoods is still a topic within local search. Consider an arbitrary neighborhood structure N , which defines for any solution s a set of neighbor solutions $N_1(s)$ as a neighborhood of depth $d = 1$. In a straightforward way, a neighborhood $N_{d+1}(s)$ of depth $d + 1$ is defined as the set $N_d(s) \cup \{s' \mid \exists s'' \in N_d(s) : s' \in N_1(s'')\}$. In general, a large d might be unreasonable, as the neighborhood size may grow exponentially. However, depths of two or three may be appropriate. Furthermore, temporarily increasing the neighborhood depth has been found to be a reasonable mechanism to overcome *basins of attraction*, e.g., when a large number of neighbors with equal quality exist [12, 13].

Large scale neighborhoods and large scale neighborhood search have become an important topic (see, e.g., [5] for a survey), especially when efficient ways are at hand for exploring them. Related research can also be found under various names; see, e.g., [92] for the idea of ejection chains.

1.2.2 Metaheuristics

The formal definition of metaheuristics is based on a variety of definitions from different authors derived from [52]. Basically, a metaheuristic is a top-level strategy that guides an underlying heuristic solving a given problem. In that sense we distinguish between a *guiding process* and an *application process*. The guiding process decides upon possible (local) moves and forwards its decision to the application process which then executes the chosen move. In addition, it provides information for the guiding process (depending on the requirements of the respective metaheuristic) like the recomputed set of possible moves.

According to [58] “metaheuristics in their modern forms are based on a variety of interpretations of what constitutes *intelligent search*,” where the term *intelligent search* has been made prominent by Pearl [91] (regarding heuristics in an artificial intelligence context; see also [118] regarding an operations research context). In that sense we may also consider the following definition: “A metaheuristic is an iterative generation process which guides a subordinate heuristic by combining intelligently different concepts for exploring and exploiting the search spaces using learning strategies to structure information in order to find efficiently near-optimal solutions.” [88].

To summarize, the following definition seems to be most appropriate: “A *metaheuristic is an iterative master process that guides and modifies the operations of subordinate heuristics to efficiently produce high quality solutions. It may manipulate a complete (or incomplete) single solution or a collection of solutions at each iteration. The subordinate heuristics may be high (or low) level procedures, or a simple local search, or just a construction method. The family of metaheuristics includes, but is not limited to, adaptive memory procedures, tabu search, ant systems, greedy randomized adaptive search, variable neighborhood search, evolutionary methods, genetic algorithms, scatter search, neural networks, simulated annealing, and their hybrids.*” [124], p. ix.

1.2.2.1 Simulated Annealing

Simulated annealing (SA) extends basic local search by allowing moves to inferior solutions [80, 35]. A basic SA algorithm may be described as follows: Iteratively, a candidate move is randomly selected; it is accepted if it leads to a solution with an improved objective function value compared to the current solution. Otherwise, the move is accepted with a probability depending on the deterioration Δ of the objective function value. The acceptance probability is computed as $e^{-\Delta/T}$, using a temperature T as control parameter. Usually, T is reduced over time for diversification at an earlier stage of the search and to intensify later.

Various authors describe a robust implementation of this general SA approach; see, e.g., [79], [77]. An interesting variant of SA is to strategically reheat the process, i.e., to perform a non-monotonic acceptance function.

Threshold accepting [37] is a modification (or simplification) of SA accepting every move that leads to a new solution that is ‘not much worse’ than the older one (i.e., it deteriorates not more than a certain threshold, which reduces with a temperature).

1.2.2.2 Tabu Search

The basic paradigm of *tabu search* (TS) is to use information about the search history to guide local search approaches to overcome local optimality (see [58] for a survey on TS). In general, this is done by a dynamic transformation of the local neighborhood. Based on some sort of memory, certain moves may be forbidden, i.e., they are set tabu. As for SA, the search may lead to performing deteriorating moves when no improving moves exist or all improving moves of the current neighborhood are set tabu. At each iteration, a best admissible neighbor may be selected. A neighbor, respectively a corresponding move, is called admissible, if it is not tabu or if an aspiration criterion is fulfilled. An aspiration criterion is a rule to eventually override a possibly unreasonable tabu status of a move. For example, a move that leads to a neighbor with a better objective function value than encountered so far should be considered as admissible.

The most commonly used TS method is based on a *recency-based* memory that stores moves, or attributes characterizing respective moves, of the recent past (*static TS*). The basic idea of such approaches is to prohibit an appropriately defined inversion of performed moves for a given period by storing attributes of the solution in a tabu list and then preventing moves that require the use of attributes in such a list.

Strict TS embodies the idea of preventing cycling to formerly traversed solutions. The goal is to provide necessity and sufficiency with respect to the idea of not revisiting any solution. Accordingly, a move is classified as tabu iff it leads to a neighbor that has already been visited during the previous search. There are two primary mechanisms to accomplish the tabu criterion: First, we may exploit logical interdependencies between the sequence of moves performed throughout the search process, as realized by, e.g., the reverse elimination method (cf., e.g., [53, 120]). Second, we may store information about all solutions visited so far. This may be carried out either exactly or, for reasons of efficiency, approximately (e.g., by using hash codes).

Reactive TS aims at the automatic adaptation of the tabu list length of static TS [15] by increasing the tabu list length when the tabu memory indicates that the search is revisiting formerly traversed solutions. A possible specification can be described as follows: Starting with a tabu list length l of 1 it is increased every time a solution has been repeated. If there has been

no repetition for some iterations, we decrease it appropriately. To accomplish the detection of a repetition of a solution, one may apply a trajectory based memory using hash codes as for strict TS.

There is a large number of additional ingredients that may make TS work well. Examples include restricting the number of neighbor solutions to be evaluated (using candidate list strategies, e.g., [97]), logical tests as well as diversification mechanisms.

1.2.2.3 Evolutionary Algorithms

Evolutionary algorithms comprise a great variety of different concepts and paradigms including genetic algorithms (see, e.g., [73, 60]), evolutionary strategies (see, e.g., [72, 106]), evolutionary programs [48], scatter search (see, e.g., [51, 54]), and memetic algorithms [87]. For surveys and references on evolutionary algorithms see also [49, 9, 85, 99].

Genetic algorithms are a class of adaptive search procedures based on principles derived from the dynamics of natural population genetics. One of the most crucial ideas for a successful implementation of a genetic algorithm (GA) is the representation of an underlying problem by a suitable scheme. A GA starts, e.g., with a randomly created initial population of artificial creatures (strings), a set of solutions. These strings in whole and in part are the base set for all subsequent populations. Information is exchanged between the strings in order to find new solutions of the underlying problem. The mechanisms of a simple GA essentially consist of copying strings and exchanging partial strings. A simple GA uses three operators which are named according to the corresponding biological mechanisms: reproduction, crossover, and mutation. Performing an operator may depend on a *fitness function* or its value (fitness), respectively. As some sort of heuristic measure, this function defines a means of measurement for the profit or the quality of the coded solution for the underlying problem and often depends on the objective function of the given problem.

GAs are closely related to *evolutionary strategies*. Whereas the mutation operator in a GA was argued to serve to protect the search from premature loss of information [60], evolutionary strategies may incorporate some sort of local search procedure (such as SD) with self adapting parameters involved in the procedure. On a simplified scale many algorithms may be classified as evolutionary once they are reduced to the following frame (see [71]). First, an initial population of individuals is generated. Second, as long as a termination criterion does not hold, perform some sort of co-operation and self-adaptation. Self-adaptation refers to the fact that individuals (solutions) evolve independently while co-operation refers to an information exchange among individuals.

Scatter search ideas established a link between early ideas from various sources—evolutionary strategies, TS and GAs. As an evolutionary approach,

scatter search originated from strategies for creating composite decision rules and surrogate constraints (see [51]). Scatter search is designed to operate on a set of points, called reference points, that constitute good solutions obtained from previous solution efforts. The approach systematically generates linear combinations of the reference points to create new points, each of which is mapped into an associated point that yields integer values for discrete variables. Scatter search contrasts with other evolutionary procedures, such as GAs, by providing unifying principles for joining solutions based on generalized path constructions in Euclidean space and by utilizing strategic designs where other approaches resort to randomization. For a very comprehensive treatment of scatter search see [81].

1.2.2.4 Cross Entropy Method

Initially proposed by [104] for the estimation of rare events, the *cross entropy* method (CE) was extended to solve combinatorial optimization problems [28]. The key ingredient in the CE is the identification of a parametric probability distribution function to be used to generate feasible solutions. Given an initial probability distribution function ϕ^0 , a converging sequence of ϕ^t is generated in such a way that each subsequent probability distribution function better captures prominent features found in high quality solutions.

At any given iteration t , ϕ^t is used to generate a population of a given cardinality. Each solution is then evaluated according to a specified merit function (or heuristic measure), e.g., the objective function value associated to each random variate, and the stochastic parameters are then updated in such a way that, in the next generation of the population, high quality solutions will have higher probabilities of being generated under the new model. The problem of updating the stochastic parameters can be solved by applying the *maximum likelihood estimator* method upon a set of “elite solutions” of the current population. In other words, given the top $\rho\%$ of the current population, the CE aims at identifying the value of the parameters of the probability distribution function that better “explains” these elite solutions. This corresponds to adjusting the model to better describe the portion of the feasible space in which good solutions have been found. The two-phase process of generation and update is repeated until convergence in probability is reached.

1.2.2.5 Ant Colony Optimization

The *ant colony optimization* (ACO) metaheuristic [32, 33] is a stochastic method based upon the definition of a construction graph and the use of a set of stochastic procedures called artificial ants. A number of frameworks for the update of stochastic parameters have been proposed. The *ant system*

is a dynamic optimization process reflecting the natural interaction between ants searching for food (see, e.g., [32, 33]). The ants' ways are influenced by two different kinds of search criteria. The first one is the local visibility of food, i.e., the attractiveness of food in each ant's neighborhood. Additionally, each ant's way through its food space is affected by the other ants' trails as indicators for possibly good directions. The intensity of trails itself is time-dependent: With time passing, parts of the trails are diminishing, meanwhile the intensity may increase by new and fresh trails. With the quantities of these trails changing dynamically, an autocatalytic optimization process is started forcing the ants' search into most promising regions. This process of interactive learning can easily be modeled for most kinds of optimization problems by using simultaneously and interactively processed search trajectories.

A comprehensive treatment of the ant system paradigm can be found in [33]. To achieve enhanced performance of the ant system it is useful to hybridize it at least with a local search component.

As pointed out by [133], there are a number of commonalities between the CE and the ACO method, especially with respect to the parameter updating rule mechanisms.

1.2.2.6 Corridor Method

The *corridor method* (CM) has been presented by [109] as a hybrid metaheuristic, linking together mathematical programming techniques with heuristic schemes. The basic idea of the CM relies on the use of an exact method over restricted portions of the solution space of a given problem. Given an optimization problem P , the basic ingredients of the method are a very large feasible space \mathcal{X} , and an exact method M that could easily solve problem P if the feasible space were not large. Since, in order to be of interest, problem P generally belongs to the class of \mathcal{NP} -hard problems, the direct application of method M to solve P becomes unpractical when dealing with real world instances, i.e., when \mathcal{X} is large.

The basic concept of a *corridor* is introduced to delimit a portion of the solution space around the incumbent solution. The optimization method will then be applied within the neighborhood defined by the corridor with the aim of finding an improved solution. Consequently, the CM defines method-based neighborhoods, in which a neighborhood is built taking into account the method M used to explore it. Given a current feasible solution $\mathbf{x} \in \mathcal{X}$, the CM builds a neighborhood of \mathbf{x} , say $\mathcal{N}(\mathbf{x})$, which can effectively be explored by employing method M . Ideally, $\mathcal{N}(\mathbf{x})$ should be exponentially large and built in such a way that it could be explored in (pseudo) polynomial time using method M .

1.2.2.7 Pilot Method

Building on a simple greedy algorithm such as, e.g., a construction heuristic the *pilot method* [38, 39] is a metaheuristic not necessarily based on a local search in combination with an improvement procedure. It primarily *looks ahead* for each possible local choice (by computing a so-called “pilot” solution), memorizing the best result, and performing the respective move. (Very similar ideas have been investigated under the acronym *rollout method* [17].) One may apply this strategy by successively performing a greedy heuristic for all possible local steps (i.e., starting with all incomplete solutions resulting from adding some not yet included element at some position to the current incomplete solution). The look ahead mechanism of the pilot method is related to increased neighborhood depths as the pilot method exploits the evaluation of neighbors at larger depths to guide the neighbor selection at depth one.

In most applications, it is reasonable to restrict the pilot process to some *evaluation depth*. That is, the method is performed up to an incomplete solution (e.g., partial assignment) based on this evaluation depth and then completed by continuing with a conventional heuristic. For a recent study applying the pilot method to several combinatorial optimization problems obtaining very good results see [123]. Additional applications can be found, e.g., in [84, 22].

1.2.2.8 Other Methods

Adaptive memory programming (AMP) coins a general approach (or even thinking) within heuristic search focusing on exploiting a collection of memory components [55, 114]. An AMP process iteratively constructs (new) solutions based on the exploitation of some memory, especially when combined with learning mechanisms supporting the collection and use of the memory. Based on the idea of initializing the memory and then iteratively generating new solutions (utilizing the given memory) while updating the memory based on the search, we may subsume various of the above described metaheuristics as AMP approaches. This also includes exploiting provisional solutions that are improved by a local search approach.

The performance as well as the efficiency of a heuristic scheme strongly depends on its ability to use AMP techniques providing flexible and variable strategies for types of problems (or special instances of a given problem type) where standard methods fail. Such AMP techniques could be, e.g., dynamic handling of operational restrictions, dynamic move selection formulas, and flexible function evaluations.

Consider, as an example, adaptive memory within TS concepts. Realizing AMP principles depends on which specific TS application is used. For example, the reverse elimination method observes logical interdependencies be-

tween moves and infers corresponding tabu restrictions, and therefore makes fuller use of AMP than simple static approaches do.

To discuss the use of AMP in intelligent agent systems, one may use the simple model of ant systems as an illustrative starting point. As ant systems are based on combining constructive criteria with information derived from the pheromone trails, this follows the AMP requirement for using flexible (dynamic) move selection rules (formulas). However, the basic ant system exhibits some structural inefficiencies when viewed from the perspective of general intelligent agent systems, as no distinction is made between successful and less successful agents, no time-dependent distinction is made, there is no explicit handling of restrictions providing protection against cycling and duplication. Furthermore, there are possible conflicts between the information held in the adaptive memory (*diverging trails*).

A natural way to solve large optimization problems is to decompose them into independent sub-problems that are solved with an appropriate procedure. However, such approaches may lead to solutions of moderate quality since the sub-problems might have been created in a somewhat arbitrary fashion. Of course, it is not easy to find an appropriate way to decompose a problem *a priori*. The basic idea of POPMUSIC is to locally optimize sub-parts of a solution, *a posteriori*, once a solution to the problem is available. These local optimizations are repeated until a local optimum is found. Therefore, POPMUSIC may be viewed as a local search working with a special, large neighborhood. While POPMUSIC has been acronymed by [112] other metaheuristics may be incorporated into the same framework, too (e.g. [107]). Similarly, in the *variable neighborhood search* (VNS) [68] the neighborhood is altered during the search in such a way that different, e.g. increasingly distant, neighborhoods of a given solution are explored. Such method can be enhanced via *decomposition*, as in the *variable neighborhood decomposition search* (VNDS) (see, e.g., [69]).

For large optimization problems, it is often possible to see the solutions as composed of parts (or chunks [129], cf. the term vocabulary building). Considering as an example the vehicle routing problem, a part may be a tour (or even a customer). Suppose that a solution can be represented as a set of parts. Moreover, some parts are more in relation with some other parts so that a corresponding heuristic measure can be defined between two parts. The central idea of POPMUSIC is to select a so-called *seed part* and a set P of parts that are mostly related with the seed part to form a sub-problem.

Then it is possible to state a local search optimization frame that consists of trying to improve all sub-problems that can be defined, until the solution does not contain a sub-problem that can be improved. In the POPMUSIC frame of [112], the set of parts P corresponds precisely to seed parts that have been used to define sub-problems that have been unsuccessfully optimized. Once P contains all the parts of the complete solution, then all sub-problems have been examined without success and the process stops.

Basically, the technique is a gradient method that starts from a given initial solution and stops in a local optimum relative to a large neighborhood structure. To summarize, both, POPMUSIC as well as AMP may serve as a general frame encompassing various other approaches.

1.2.3 *Miscellaneous*

Target analysis may be viewed as a general learning approach. Given a problem, we first explore a set of sample instances and an extensive effort is made to obtain a solution which is optimal or close to optimality. The best solutions obtained provide *targets* to be sought within the next part of the approach. For instance, a TS algorithm may be used to bias the search trajectory toward already known solutions (or as close to them as possible). This may give some information on how to choose parameters for other problem instances.

A different acronym in this context is *path relinking* (PR) which provides a useful means of intensification and diversification. Here new solutions are generated by exploring search trajectories that combine elite solutions, i.e., solutions that have proved to be better than others throughout the search. For references on target analysis and PR see, e.g., [58].

Considering local search based on *data perturbation*, the acronym *noising method* may be related to the following approach, too. Given an initial feasible solution, the method performs some data perturbation [111] in order to change the values taken by the objective function of a problem to be solved. On the perturbed data a local search may be performed (e.g., following a SD approach). The amount of data perturbation (the noise added) is successively reduced until it reaches zero. The noising method is applied, e.g., in [23] for the clique partitioning problem.

The key issue in designing *parallel algorithms* is to decompose the execution of the various ingredients of a procedure into processes executable by parallel processors. Opposite to ant systems or GAs, metaheuristics like TS or SA, at first glance, have an intrinsic sequential nature due to the idea of performing the neighborhood search from one solution to the next. However, some effort has been undertaken to define templates for parallel local search (see, e.g., [119, 117, 26, 116]). A comprehensive treatment with successful applications is provided in [6]. The discussion of parallel metaheuristics has also led to interesting hybrids such as the combination of a population of individual processes, agents, in a cooperative and competitive nature (see, e.g., the discussion of *memetic algorithms* in [87]) with TS.

Neural networks may be considered as metaheuristics, although we have not considered them here; see, e.g., [108] for a comprehensive survey on these techniques for combinatorial optimization. On the contrary, one may use metaheuristics to speed up the learning process regarding artificial neural networks; see [7] for a comprehensive consideration.

Require: \mathbf{x}^k incumbent solution; Ω^k current set of heuristic parameters
Ensure: \mathbf{x}^{k+1} next solution

1. $\mathcal{N}(\mathbf{x}^k) \leftarrow \text{neighborhood_definition}(\mathbf{x}^k, \Omega^k)$
2. $\mathbf{x}^{k+1} \leftarrow \text{neighborhood_exploration}(\mathcal{N}(\mathbf{x}^k))$
3. $\Omega^{k+1} \leftarrow \text{parameters_update}()$

Fig. 1.1 General Metaheuristic Iteration.

Stochastic local search (SLS) is pretty much all we know about local search but enhanced by randomizing choices. That is, an SLS algorithm is a local search algorithm making use of randomized choices in generating or selecting candidate solutions for given instances of optimization problems. Randomness may be used for search initialization as well as the computation of search steps. A comprehensive treatment of SLS is given in [75].

Furthermore, recent efforts on problems with multiple objectives and corresponding metaheuristic approaches can be found in [78, 41]. See, e.g., [105] for some ideas regarding GAs and fuzzy multi-objective optimization.

1.3 A Taxonomy

In this section, we present a taxonomy of metaheuristics along a single dimension of analysis. The driving factor is the way in which the neighborhood is defined with respect to each metaheuristic approach. Alternative classifications have been proposed, e.g., in [19, 64].

From a general perspective, each metaheuristic paradigm can be seen as made up by three major ingredients, which are repeatedly used at each iteration until specified stopping criteria are reached. A generalization of a single iteration of a metaheuristic scheme is given in Figure 1.1.

As illustrated in Step 1 of Figure 1.1, a common ingredient of each metaheuristic paradigm is the existence of a rule aimed at iteratively guiding the search trajectory, i.e., a set of rules to define a neighborhood. In turn, such neighborhood demarcates which solutions can be reached starting from the incumbent solution. In line with this observation, a possible dimension along which a taxonomy of metaheuristics can be created is given by the way in which neighborhoods are built. A classification of metaheuristics along this dimension leads to the definition of at least two broad classes:

- *model-based heuristics*: as in [133], with this term we refer to metaheuristic schemes where new solutions are generated by using a model. Consequently, the neighborhood is implicitly defined by a set of parameters, and iteratively updated during the search process;

- *method-based heuristics*: as in [109], with this term we make reference to heuristic paradigms in which new solutions are sought in a neighborhood whose structure is dictated by the method used to explore the neighborhood itself. Consequently, the neighborhood is implicitly defined by the predetermined method employed.

By observing the underlying philosophy of these two broad classes, a clear dichotomy arises: on the one hand, model-based heuristics tackle the original optimization problem by defining and iteratively updating a model aimed at identifying prominent features in good solutions and at replicating these features in future solutions. Consequently, what determines whether a point belongs to the neighborhood is the set of parameters that defines the model and the ‘probability’ of generating such point under the current model. On the other hand, method-based heuristics are driven by the technique used to solve a “reduced” version of the original problem, i.e., a problem in which only a subset of the original solution space is considered. Consequently, in method-based heuristics what dictates the structure and shape of the neighborhood is the optimization method employed to explore the neighborhood itself, whether it be a classical mathematical programming technique, e.g., branch and bound, dynamic programming, etc., or a simple enumeration-based technique.

Model-based heuristics are generally based upon the identification of a set of parameters, defining a model that, in turn, well captures some features of the search space. This type of heuristics heavily relies on a set of update schemes, used to progressively modify the model itself in such a way that, after each update, the possibility of obtaining higher quality solutions under the new model is increased. Consequently, in Step 1 of the general `Metaheuristic_Iteration()`, all the solutions that “comply” with the requirements enforced by the model upon the search space are included in the current neighborhood. A special role is played by Step 3 of the same algorithm, in which the parameters of the model are updated via the application of learning mechanisms. In this phase, modifications are applied to the model and/or its parameters to reflect insight collected and generated during the search phase.

A well-known paradigm that can be interpreted under the philosophy of the model-based method is the CE, where a stochastic model is continually updated to reflect the findings of the last iteration of the search process. Other metaheuristics that can be seen as model-based are ACO (as well as other methods belonging to the *Swarm Intelligence* field), where a construction graph and stochastic procedures called *ants* are employed; semi-greedy heuristics [70], including the greedy randomized adaptive search procedure *GRASP* [43], where the greedy function that guides the selection of the best candidates defines a stochastic model; and GAs, where adaptive search procedures based upon genetics are put into place. The GA paradigm heavily relies on a model, defined by a set of operators, that determines which solutions will be included in the next generation, i.e., in the current neighborhood.

More generally, evolutionary algorithms could similarly be included into the category of model-based metaheuristics.

On the other side of the spectrum we find metaheuristic paradigms driven by a method, rather than by a model. The basic ingredient of such an approach is the existence of a search method, whether it be an exact method or a heuristic method, that is used to explore the neighborhood itself. Consequently, the size and cardinality of the neighborhood depend upon the ability of the method itself to explore the portion of the search space included in the neighborhood. Within the class of method-based metaheuristics we can introduce a further level of classification, according to the nature of the method employed to explore the neighborhood. Broadly speaking, method-based heuristics could be divided into two categories, those for which classical mathematical programming techniques are employed to explore the neighborhood and those for which enumeration-based techniques are used to conduct the exploration of the basin of solutions.

A cardinal concept for the method-based heuristics is connected to the introduction of a “distance” metric. A distance is used to draw the boundaries of the neighborhood around the incumbent solution, in such a way that only points whose distance from the incumbent solution is within a threshold are included in the neighborhood itself. Consequently, the neighborhood is explicitly defined by the notion of distance adopted. On the other hand, the definition of the threshold distance is strongly connected with the capabilities of the method used to explore the neighborhood. In other words, the cardinality of the neighborhood is chosen in such a way that, on the one hand, it will be large enough to have a reasonable chance of containing a solution better than the incumbent one and, on the other hand, small enough to be explored by employing the method at hand in a reasonable amount of computational time.

Historically, the first metaheuristics developed might be regarded as belonging to this class, e.g., SA or TS. Let us consider, e.g., the TS metaheuristic. Given an incumbent solution \mathbf{x}^i , a distance function $d(\mathbf{x}_1, \mathbf{x}_2)$ and a threshold value δ , only solutions for which $d(\mathbf{x}^i, \mathbf{x}) \leq \delta$ will be included into the current neighborhood. Once the neighborhood definition phase is terminated, a method capable of exploring such neighborhood in a reasonable amount of computational time is employed to find a possibly better solution. Consequently, while the neighborhood is explicitly defined by the value of δ , it is possible to say that such neighborhood is implicitly defined by the method used, since the value of δ depends upon the capabilities of the method itself.

It is worth noting that these metaheuristics can, and in general, do use a set of parameters to refine the definition of the neighborhood. For example, let us once more consider the case of the TS metaheuristic. The neighborhood is first defined according to a “distance” from the incumbent solution and, then, refined via the application of, e.g., the tabu list, with the effect of eliminating some of the solutions from the current neighborhood. However, it should be

evident that the major ingredient used in determining the neighborhood is still related to the concept of distance from the incumbent solution.

Other metaheuristics that fit into this category are VNS [68] and the *pilot method*. In VNS, the neighborhood is altered during the search in such a way that increasingly distant neighborhoods of a given solution are explored. However, a “method”, e.g., the local search routine, must be applied to perform the search over the neighborhood. Thus, the way in which neighborhoods are built is influenced by the method used to explore the portion of the search space at hand. Similarly, in the pilot method, the core idea is that the neighborhood is defined by a look-ahead mechanism. Consequently, there is a method (even a simple local search) that determines the shape, or the deepness, of the neighborhood itself.

More recently, metaheuristics employing classical mathematical programming techniques to explore the neighborhood have been proposed. Let us consider the case of the CM. After defining a corridor around the incumbent solution, an optimization method is then applied within the neighborhood defined by the corridor with the aim of finding an improved solution. Consequently, in Step 1 of the general metaheuristic iteration of Figure 1.1, only solutions within the corridor, or within a predefined distance from the incumbent, will be included in the neighborhood. As previously mentioned, an optional feature of the `neighborhood_definition()` phase is the application of a set of criteria to refine the neighborhood, e.g., a tabu list, aspiration criteria, etc. Step 2 of the algorithm relies on the use of either an enumeration-based technique, or a classical mathematical programming technique (branch and bound, dynamic programming, etc.) to explore the neighborhood. Finally, Step 3 of the algorithm, `parameters_update()` can include, e.g., the dynamic update of distance and corridor parameters, depending upon the current status of the search.

Beside the corridor method, other concepts that fall into this category are, e.g., constraint programming [103], in which the corridor is defined by the subsequent application of constraints and conditions to be satisfied by the solution, and local branching [47], in which the neighborhood is defined by introducing linear inequalities, called local branching cuts. Once a restricted neighborhood is so defined, an exact technique, i.e., linear programming, is used in the spirit of the branch and bound framework.

A further example which might be interpreted as a method-based technique is the POPMUSIC framework [112], where one wants to solve, preferably to optimality, smaller portions of the solution space, based upon an available feasible solution.

More recently, the *relaxation induced neighborhood search* method (RINS) has been introduced [27]. The RINS defines a neighborhood exploiting information contained in the linear programming (LP) relaxation and can naturally be seen as a method-based framework for mixed integer programs (MIP). The central idea of the method is related to the exploitation of a “relaxed” solution to define a core problem, smaller than the original MIP. The core

problem is identified at any given node of the branch and cut tree, by first fixing all variables that have the same values both in the incumbent (feasible) solution and the relaxed solution. Next, a branch-and-cut framework is used to solve to optimality the reduced problem on the remaining variables, called sub-MIP. The process is iteratively applied at any node of the global tree, since the LP relaxation induced solution is different and, therefore, gives raise to different sub-MIPs.

1.4 Hybrids with Exact Methods

In recent years, a lot of attention has been devoted to the integration, or hybridization, of metaheuristics with exact methods (see, e.g., [95, 96] for a survey and a taxonomy about hybrid approaches in combinatorial optimization, respectively.) In this section, we use the term *hybrid* in a somehow restrictive way, since we classify as *hybrid approaches* only those approaches that combine the use of exact techniques with metaheuristic frameworks. Consequently, algorithms that combine together different metaheuristics are not included in this analysis although they could be and are also termed hybrid.

This exposition also relates to the term *Matheuristics*, which describes works that also are along these lines, e.g., exploiting mathematical programming techniques in (meta)heuristic frameworks or on granting to mathematical programming approaches the cross-problem robustness and constrained-CPU-time effectiveness which characterize metaheuristics. Discriminating landmark is some form of exploitation of the mathematical formulation of the problems of interest [67].

Generally speaking, hybrid algorithms present a so-called “master-slave” structure of a guiding process and an application process. Either (i) the metaheuristic acts at a higher level and controls the calls to the exact approach, or (ii) the exact technique acts as the master and calls and controls the use of the metaheuristic scheme.

Hybrid algorithms of type (i) are such that the *definition* of the neighborhood follows the logic of a metaheuristic, while the *exploration* of the neighborhood itself is left to the exact approach. From this perspective, the metaheuristic acts as the master by defining the size and boundaries of the neighborhood and by controlling repeated calls to the exact method, which, in turn, acts as an application process, by exploring each neighborhood in an exact fashion. Algorithms that fall into this category are, e.g., those inspired by the CM, in which large scale neighborhoods are searched exhaustively through an exact method applied on a sub-portion of the search space. The call to the exact method is managed by a scheme that heuristically defines a corridor around an incumbent solution (cf. the previous section).

A similar philosophy is shared by the *large scale neighborhood search* (see, e.g., [5] for a survey), in which exponentially large neighborhoods are searched to optimality by means of, e.g., ad-hoc enumeration schemes, dynamic programming schemes, etc.

Along the same line, the RINS as well as local branching could be seen as algorithms of type (i), at least in *spirit*. For example, even though the RINS is casted into a branch and cut approach, and, therefore, the guiding process is an exact approach, the logic of the method is centered upon metaheuristic-type features, such as neighborhood definition, diversification and intensification. It is worth noting, though, that in these two approaches, no real metaheuristic is ever deployed, since they entirely rely on the branch and bound framework. However, they can still be seen as hybrid approaches because of the embedded metaheuristic philosophy that drives the search process.

On the other hand, we also have hybrid approaches of type (ii), in which the metaheuristic scheme is embedded into the solver. Modern branch and cut solvers exploit the potentials of (meta)heuristics to quickly get good quality solutions, especially at early stages of the tree exploration. Related bounds are then employed to prune branches of the tree and, consequently, contribute to speed up the search process and to reduce the overall computational effort. Since it has been observed that in some important practical cases MIP solvers spend a large amount of computational time before finding the first feasible solution, [46] introduced a heuristic scheme, further improved in [3, 16], called the *feasibility pump*, aimed at quickly finding good quality initial solutions. Such initial solutions are obtained via a sequence of roundings, based upon continuous relaxation solutions, that converge to a feasible MIP solution. Clearly, such heuristic-type schemes can also be used to quickly find initial solutions to be fed to type (i) hybrid algorithms, such as the CM, the RINS as well as local branching (see also [56, 57, 40] for heuristic methods for MIP feasible solution generation).

In a fashion similar to hybrid approaches of type (ii), some researchers have also employed metaheuristic schemes for column generation and cut generation within branch and price and branch and cut frameworks, respectively (see, e.g., [44] and [94]). In addition, one may investigate hybrids of branch and bound and metaheuristics, e.g., for deciding upon branching variables or search paths to be followed within a branch and bound tree (see, e.g., [130] for an application of reactive TS). Here we may also use the term *cooperative solver*.

A key question that arises when designing a hybrid algorithm concerns which components should be “hybridized” to create an effective algorithm. While providing an all-encompassing rule for hybridization does not seem to be a feasible approach, from the analysis of the state of the art of hybrid algorithms some interesting guidelines emerge.

A method-based approach is centered upon the exploitation of an effective “method” to solve the problem at hand. Consequently, the starting point lies

in the identification of the most effective(s) method(s) with respect to the optimization problem. For example, the design of a CM inspired algorithm requires previous knowledge about which method could effectively tackle the problem if this were of reduced size. Thus, the identification of the method to be used constitutes the central point in the design of the algorithm itself.

A basic ingredient of a method-based algorithm concerns the heuristic rule used to draw the boundaries of the neighborhood upon which the method will be applied, which is the *design of the neighborhood* itself in terms of how large the neighborhood should be. Size and boundaries of such neighborhood depend on the “power” of the method used, i.e., on its ability to explore large portions of the solution space in a reasonable amount of computational time. While determining the appropriate dimension of the neighborhood for the method at hand is an issue of algorithm fine tuning (as presented in Section 1.6), some general considerations are related to the fitness landscape analysis as well as the computational complexity of the method itself. Roughly speaking, given an optimization method and its worst case computational complexity in terms of size of the input, it is possible to determine the maximum size of the neighborhood that guarantees running times below a desired threshold. On the other hand, since computational complexity analysis mainly deals with worst case scenarios, it seems beneficial to employ fitness landscape analysis techniques (e.g., connectivity measures) to draw tightest complexity bounds that translate directly into larger neighborhoods.

Another guideline is provided by the *intensification-diversification trade-off*. By reviewing hybrid algorithms proposed in the literature, many times it is possible to identify a predominant focus, in the sense that some algorithms put a higher emphasis on diversification of solutions, while others emphasize the intensification of the search in promising regions. For example, as illustrated in [47], the application of valid inequalities in the spirit of local branching fosters the intensification of the search in a given neighborhood, hence allowing to find good quality solutions early on in the search process. In a similar fashion, the CM seems to put more emphasis on intensifying the search within a promising region, without defining specific restarting mechanisms to achieve diversification. On the other hand, a method such as the RINS, based upon a solution of the LP relaxation of the MIP, puts more emphasis on diversification, since at each node of the search tree a different LP induced solution is used and, consequently, different solutions feasible with respect to the MIP will be produced.

Finally, an interesting line of research aimed at grasping a clearer understanding of why some search techniques are successful on a given problem class is related to the *fitness landscape analysis*. As mentioned in [18, 115], a central measure of landscape structure is the fitness-distance correlation, which captures the correlation between objective function value and the length of a path to an optimal solution within the fitness landscape. Such a measure is used to explain why local search techniques perform well in tackling certain problems. However, the link between problem difficulty and

fitness landscape is, as of today, not completely understood (see also the idea of target analysis mentioned above).

Fitness landscape analysis can be used with at least two goals in mind:

- on the one hand, as brought out in [126], this kind of analysis helps to understand what makes a problem hard or, conversely, well suited, for a specific search technique. Information such as fitness-distance correlation, ruggedness, nodes connectivity, and drifting can be exploited to design an effective metaheuristic scheme as well as to identify which components should be hybridized;
- on the other hand, as illustrated in [24], fitness landscape analysis can help to identify which formulation of the same problem will be more suitable with respect to an available algorithm. For example, [24] were able to “predict” the behavior of a VNS algorithm upon two different formulations of the Golomb Ruler problem and, consequently, to select the formulation that better fitted with the potentials of their algorithm.

While this field of study seems promising in grasping a better understanding of the “hows” and “whys” of metaheuristics, an important issue of generalization of results has already been noticed. As mentioned in [126], the results obtained so far have mainly been used *a posteriori*, to justify the use of a given algorithm and its features. However, it is unclear how this kind of analysis can be extended to develop improved algorithms, since there is no clear understanding about the general validity of the findings of the proposed models. However, it is worth noting that, from the methodological perspective, the contribution of the fitness landscape analysis is far from negligible, since its focus is well oriented toward interpreting and partially explaining successes and failures of metaheuristic-based algorithms.

1.5 General Frames: A Pool-Template

An important avenue of metaheuristics research refers to general frames (e.g., to explain the behavior and the relationship between various methods) as well as the development of software systems incorporating metaheuristics (eventually in combination with other methods). Besides other aspects, this takes into consideration that in metaheuristics it has very often been appropriate to incorporate a certain means of diversification versus intensification to lead the search into new regions of the search space. This requires a meaningful mechanism to detect situations when the search might be trapped in a certain area of the solution space. Therefore, within intelligent search the exploration of memory plays a most important role.

In [64] a *pool template* (PT) is proposed as can be seen in Figure 1.2. The following notation is used. A pool of $p \geq 1$ solutions is denoted by P . Its input and output transfer is managed by two functions which are called IF

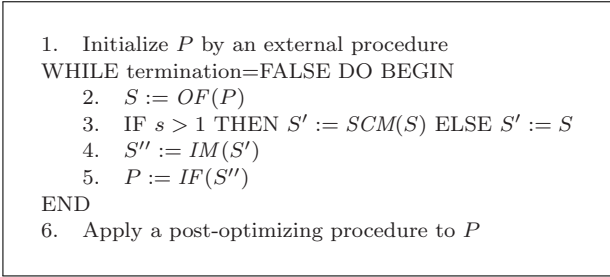


Fig. 1.2 Pool Template.

and OF , respectively. S is a set of solutions with cardinality $s \geq 1$. A solution combination method (procedure SCM) constructs a solution from a given set S , and IM is an improvement method.

Depending on the method used, in Step 1 either a pool is completely (or partially) built by a (randomized) diversification generator or filled with a single solution which has been provided, e.g., by a simple greedy approach. Note that a crucial parameter that deserves careful elaboration is the cardinality p of the pool. The main loop, executed until a termination criterion holds, consists of Steps 2–5. Step 2 is the call of the output function which selects a set of solutions, S , from the pool. Depending on the kind of method represented in the PT, these solutions may be assembled (Step 3) to a (set of) working solution(s) S' which is the starting point for the improvement phase of Step 4. The outcome of the improvement phase, S'' , is then evaluated by means of the input function which possibly feeds the new solution into the pool. Note that a post-optimizing procedure in Step 6 is for facultative use. It may be a straightforward greedy improvement procedure if used for single-solution heuristics or a pool method on its own. As an example we quote a *sequential* pool method, the TS with PR in [11]. Here a PR phase is added *after* the pool has been initialized by a TS. A *parallel* pool method on the other hand uses a pool of solutions *while* it is constructed by the guiding process (e.g., a GA or scatter search).

Several heuristic and metaheuristic paradigms, whether they are obviously pool-oriented or not, can be summarized under the common PT frame. We provide the following examples:

- a) Local Search/SD: PT with $p = s = 1$.
- b) SA: $p = 2, s = 1$ incorporating its probabilistic acceptance criterion in IM . (It should be noted that $p = 2$ and $s = 1$ seems to be unusual at first glance. For SA we always have a current solution in the pool for which one or more neighbors are evaluated and eventually a neighbor is found which replaces the current solution. Furthermore, at all iterations throughout the search the so far best solution is stored, too (even if no real interaction between those two stored solutions takes place). The same is also valid for

- a simple TS. As for local search the current solution corresponds to the best solution of the specific search, we have $p = 1$.)
- c) Standard TS: $p = 2$, $s = 1$ incorporating adaptive memory in *IM*.
 - d) GAs: $p > 1$ and $s > 1$ with population mechanism (crossover, reproduction and mutation) in *SCM* of Step 3 and without the use of Step 4. (One might argue that $p \geq 1$ is also possible.)
 - e) Scatter Search: $p > 1$ and $s > 1$ with subset generation in *OF* of Step 2, linear combination of elite solutions by means of the *SCM* in Step 3, e.g., a TS for procedure *IM* and a reference set update method in *IF* of Step 5.
 - f) PR (as a parallel pool method): $p > 1$ and $s = 2$ with a PR neighborhood in the *SCM*. Facultative use of Step 4.
 - g) CE: $p > 1$ and $s = \rho p$, with $\rho \in (0, 1)$, where the *SCM* is used to update the underlying stochastic model by capturing features of solutions in *S* and without the use of Step 4.

1.6 Fine Tuning and Evaluation of Algorithms

In this section, we discuss two important issues related to the design and analysis of metaheuristics and related algorithms. On the one hand, as mentioned in the introduction, one key factor that has a bearing on the overall performance of most of these algorithms is the calibration of the algorithmic parameters. Thus, a natural question is about how to select an appropriate set of values for these parameters. This important topic in metaheuristic design goes under the name of *fine tuning* of algorithmic parameters.

A second relevant issue, related to the analysis of an algorithm, is concerned with the empirical evaluation of the performance of the algorithm itself. Due to the stochastic nature of many metaheuristic schemes, a problem of reproducibility of results arises. Many researchers advocate the definition of a set of standards to increase the objectivity of the “ranking” and evaluation of metaheuristics. However, the literature does not seem to be mature with respect to this topic.

1.6.1 Fine Tuning of Metaheuristics

According to [4], there is evidence that 10% of the time required to develop a new metaheuristic is devoted to the actual development and that the remaining 90% is spent on fine tuning of algorithmic parameters. In addition, fine tuning of parameters strongly affects the final performance of an algorithm. For these reasons, it is of paramount importance to make a concerted effort in identifying and establishing a set of “standard” techniques to fine-tune a metaheuristic. One of the major achievements of such an effort would be to

offset parameter specific issues in evaluating an algorithm. In addition, reproducibility of results would also be enhanced by such an approach, by making transparent the way in which parameter values should be set to tackle a given problem instance.

In the literature, some attempts to use statistically robust methods have been presented. For example, in [4], a tool called CALIBRA is proposed as a procedure that finds good values for up to five algorithmic parameters. They exploit Taguchi fractional factorial design to reduce the overall number of trials required to train the model. Interestingly, the authors tested CALIBRA on a set of six different algorithms and use hypothesis testing to assert whether the parameter values suggested by the tool allow to find solutions which are significantly better than those proposed by the original authors of the algorithms.

In [25], a four-step procedure is proposed, in which a two-level factorial design is coupled with linear regression to find a linear approximation of the response surface of the set of parameters. Subsequently, a gradient descent technique is employed to find a “good” value for each parameter. Finally, the method is applied on a set of benchmark problems, and results are collected to show that a good parameter setting leads to improvements in the objective function value. A similar technique has been employed in [21], where a circumscribed central composite design is used to generate observations and a higher degree polynomial is then used to approximate the response surface of a set of parameters. Once the response surface is known, a global optimization method is employed to find the optimal parameter value with respect to the surface. The model is built on a set of training instances and then validated on a different set of testing instances of some lot sizing problems.

In [89], a nonlinear response surface is used to capture the impact of parameters on a SA algorithm. A good set of values of such parameters is then determined via a modified simplex method for nonlinear programming, that can be used to deal with bounds on parameter values. The proposed method is tested on three different combinatorial optimization problems and the collected results are then compared with those of other SA implementations, whose fine tuning was achieved via extensive experiments. The results show that there is no statistical difference in performance between the two types of SA algorithms and, consequently, that the proposed approach can be used to conduct an “automatic” fine tuning. A similar approach was followed in [90], where a full-factorial design was employed to define a response surface for up to four parameters of a GA. The fine-tuned algorithm was finally tested on a pool of DNA sequence assembly problems and results were collected to illustrate the effectiveness of the procedure.

A somehow different approach is proposed in [131], where two statistical tests are employed to fine-tune a set of algorithmic parameters of a TS algorithm and to validate the proposed statistical model. First, a Friedman test is used to detect significance of a specific parameter upon the algorithm performance and, via a series of pairwise comparisons, to identify a good pa-

parameter value; subsequently, a Wilcoxon test is employed to verify whether there is a statistically significant difference between any two F-runs of the algorithm. The TS algorithm is then used upon a set of problems drawn from the telecommunication network design field.

A relevant issue in fine tuning methods is related to the “budget” available to carry on the experiments, in the sense that the type of techniques used depends on the amount of time and computational power available. These elements affect the number of factors (parameters) and the number of levels (values) that can be considered in the experimental design. In this regard, as illustrated in [4], the use of a “fractional factorial design” can help in reducing the number of runs required to collect results.

A second issue concerns the analysis of results in terms of robustness and sensitivity. In many circumstances, it is necessary to provide not only a “good” set of parameter values but also a measure of robustness and sensitivity with respect to those parameters. For this reason, the experiment should be designed in such a way that training and testing sets are used to create and to validate the model, respectively.

Finally, an important issue that needs to be addressed when fine tuning algorithmic parameters is the ultimate goal that one wants to achieve. In the literature, many times it is implicitly assumed that the goal is to maximize solution quality and, with this in mind, one wants to find a good set of parameters. However, alternative goals could also be desirable, such as, e.g., minimizing computational time to a target solution, generating a pool of solutions with maximum “diversification” [63], obtaining good quality solutions early on in the search process, as well as a weighted combination of such goals. Consequently, the fine tuning process should be designed in such a way that a trade-off between conflicting goals could be achieved.

In the spirit of what we have seen in the reactive TS fine tuning also concerns means of autoadaptivity of parameter settings. In a more general setting an excellent recent treatment is provided in [14].

1.6.2 Empirical Evaluation of Metaheuristics

As stated in [10], it is important to “promote thoughtful, well-planned, and extensive testing of heuristics, full disclosure of experimental conditions, and integrity in and reproducibility of the reported results.”

With this aim in mind, a number of interesting issues have been proposed with respect to how to measure the performance of a metaheuristic. For example, some relevant issues connected with the evaluation of metaheuristics are:

- computational experiment design, with a clear objective in mind (e.g., to illustrate solution quality performance, robustness, quality versus time, etc.);

- testing on benchmark instances, if available, or with respect to (dual) gaps if possible. Results should be reported with measures of variability and robustness;
- if no comparison with other algorithms from the literature is possible, the algorithm should at least be compared with a simple random restart procedure;
- identification of the contribution of each major feature of the algorithm on the final performance, to detect which features are relevant in achieving the declared level of performance;
- measurement of statistically significant differences among algorithms, with the aim of *ranking* a pool of algorithms designed to tackle the same class of problems.

In this section, we focus on two major issues: *reproducibility* and *ranking* of an algorithm. On the one hand, the experimental results should be reported in such a way that reproducibility is ensured while, on the other hand, a statistically sound comparison of the proposed method with other approaches from the literature should be presented with the aim of detecting meaningful improvements in the state of the art.

The issue of reproducibility has been addressed in [10]. In order to be of any scientific value, the experiments used to assess an algorithm in general, and a metaheuristic in particular, should be entirely reproducible by others. On the one hand, documentation will help in fostering reproducibility. For this reason, not only the algorithmic steps but also specifics about the implementation, the data structure employed, the parameter settings, the random number process generation (if applicable), etc., should be provided. In addition, making available to the community the source code as well as the instances used fosters reproducibility and enhances the quality of the scientific work.

Many statistical tests have been proposed to determine whether one algorithm outperforms another. The way in which this issue should be addressed is influenced by a number of relevant factors, such as, e.g., the nature of the algorithm itself (deterministic or stochastic), the size of the available sample (test bed), and, of course, the measure(s) used to draw conclusions about the quality of a given algorithm. In [31], a taxonomy of statistical questions, applied to learning algorithms, is presented and some hints about how to address the issue of comparing algorithms with respect to each category are given.

In the operations research field, the issue of comparing two algorithms with each other arises in at least two different contexts:

- on the one hand, one might want to compare two different versions of the same algorithm using different parameter settings to *fine tune* the algorithm itself. In this case, the practitioner aims at detecting whether a given set of parameters produces better performance;

- on the other hand, to *validate* an algorithm, a researcher will have to compare the performance of the proposed algorithm against, at least, those of the best algorithm available for a specific class of problems. In this case, the underlying assumption, quite common in the operations research field, is that the algorithm is designed to perform in a specific domain, i.e., to tackle a single class of problems, and that it is possible to identify a state of the art algorithm.

While the literature on the use of statistical analysis for hypothesis testing is abundant (e.g., [132], [50]), in general, only rudimentary techniques are used to present results and assert the quality and soundness of the same results produced by a new metaheuristic. Other fields seem more mature when it comes to statistical validation of results. For example, the machine learning community has become increasingly aware of the importance of validating results with sound statistical analysis [29]. In [31], five statistical tests for comparison of two different algorithms are presented. These tests are experimentally compared with respect to the probability of incurring Type I errors (incorrectly detecting a difference when no difference exists). More sophisticated tests are proposed in [29], which can be used to compare multiple algorithms over multiple data sets.

In the field of metaheuristics, the use of statistical tests is somehow present when fine tuning algorithms. For example, as mentioned in Section 1.6.1, [4], [131], [89], use different statistical tests to assert the quality of the fine tuning technique. However, the literature concerning authors that employ statistical analysis to *compare* two, or a set, of metaheuristics upon a given set of instances is quite scanty. A good introduction to the topic can be found in [113]. In this paper, only the case in which two algorithms are to be compared is presented. A number of statistical tests is reviewed, both parametric and non-parametric tests, and a new non-parametric test is proposed (see also [110] for some ideas).

The major issue in identifying which test should be used to compare two algorithms is the identification of the key factor to be used to “judge” such algorithms. For example, in the contest of optimization, a common measure of quality is the objective function value. Consequently, one might want to compare two algorithms, say algorithm A and algorithm B, in terms of a set of related objective function values, say z^A and z^B . The null hypothesis is then defined as $H_0 : z^A - z^B = 0$ and a one-sided or a two-sided test is designed. In general, classical parametric approaches are used to test such hypothesis such as, e.g., a (paired) t-test, which checks whether the average difference in performance over the data set is significantly different from zero. However, as brought out by many authors (see, e.g., [29]), the t-test suffers from a number of weaknesses, mainly that the underlying assumption of normal distributions of the differences between the two random variables could not be realistic, especially when the number of benchmark instances available is not large.

An alternative approach relies on the use of non-parametric tests for comparing proportions. Let us suppose that we count a success for algorithm A

every time algorithm A outperforms algorithm B on a given instance (e.g., in terms of objective function value, running time, etc.). The researcher is interested in estimating the success probabilities of the two algorithms, say p_A and p_B . After estimating empirically such probabilities, a non-parametric test (e.g., McNemar test, Fisher exact test) could be used to “rank” such algorithms in terms of effectiveness, indicating which algorithm is more successful on the same data set. A non-parametric test that can be used to “rank” algorithms based upon the success probabilities is proposed by [113]. In addition, such test overcomes one important limitation of the McNemar test, namely the fact that these tests require pairwise comparisons. Many times in the field of metaheuristics researchers test their algorithms on randomly generated instances, whose generation process is described while the actual instances are not made available to the community. Hence, it is not always possible to compare two algorithms on the same set of instances and the McNemar test might not be significant. The test proposed in [113] overcomes such obstacle by taking into account the total number of runs of algorithm A and of algorithm B (which might be different) in comparing proportions and proves to be more powerful than the Fisher test.

More complicated scenarios could be envisioned when more than one dimension of evaluation is taken into account. For example, quite often a trade off between solution quality and computational time arises. Consequently, one might be interested in evaluating, or ranking, algorithms along these two conflicting dimensions. In addition, if the metaheuristic is part of a broader hybrid algorithm, such as those mentioned in Section 1.4, a valid measure of effectiveness could be the “degree of diversification” of the solutions provided by the metaheuristic. Thus, one might want to evaluate two algorithms according to their ability of generating diversified solutions, which are, then, fed to a global optimizer. In these cases, where one wants to evaluate the effect of a pool of factors, a multi-objective problem could be defined, in such a way that a “weight” is assigned to each criterion and the “success” of an algorithm is measured as the weighted sum of the different criteria.

Finally, one last word could be spent about which algorithm(s) the proposed metaheuristic should be compared against. Generally speaking, the goal is to compare the new metaheuristic with established techniques, e.g., the best algorithm available in the literature. However, as pointed out in [10], rather than reporting comparisons with results produced on different machines and different instances, it is better to obtain or, in its defect, to recode, existing algorithms and conduct a fair comparison of the two algorithms on the same set of instances and the same machine. On the other hand, if no algorithms have been proposed for the problem at hand, a more general method, e.g., one based on linear programming or integer programming, could be used to obtain bounds on the objective function values. Whenever a stochastic scheme is proposed, at least a comparison with a simple random restart procedure should be carried on, to show that the proposed algorithm

performs *significantly better* (in a statistical sense) than the random restart procedure.

1.7 Optimization Software Libraries

Besides some well-known approaches for reusable software in the field of exact optimization (e.g., CPLEX² or ABACUS³) some ready-to-use and well-documented component libraries in the field of local search based heuristics and metaheuristics have been developed; see the contributions in [125].

The most successful approaches documented in the literature are the Heuristic OpTimization FRAMEwork HOTFRAME of [45] and EASYLOCAL++ of [30]. HOTFRAME, as an example, is implemented in C++, which provides adaptable components incorporating different metaheuristics and an architectural description of the collaboration among these components and problem-specific complements. Typical application-specific concepts are treated as objects or classes: problems, solutions, neighbors, solution attributes and move attributes. On the other side, metaheuristic concepts such as different methods described above and their building-blocks such as tabu criteria or diversification strategies are also treated as objects. HOTFRAME uses genericity as the primary mechanism to make these objects adaptable. That is, common behavior of metaheuristics is factored out and grouped in generic classes, applying static type variation. Metaheuristics template classes are parameterized by aspects such as solution spaces and neighborhood structures.

Another well-known optimization library is the COIN-OR library⁴, an open-source suite for the optimization community. An effort in the development of standards and interfaces for the interoperability of software components has been put forth. Some classes that implement basic ingredients of metaheuristics, e.g., Tabu Search Project, have been developed. However, the development of general software frameworks for metaheuristic paradigms, although strategic in defining standards and commonalities among approaches, is still in its infancy.

1.8 Conclusions

Over the last decades metaheuristics have become a substantial part of the optimization stockroom with various applications in science and, even more

² www.ilog.com

³ www.informatik.uni-koeln.de/abacus

⁴ <http://www.coin-or.org>

important, in practice. Metaheuristics have become part of textbooks, e.g. in operations research, and a wealth of monographs (see, e.g., [118, 58, 86, 36]) is available. Most important in our view are general frames. Adaptive memory programming, Stochastic Local Search, an intelligent interplay of intensification and diversification (such as ideas from POPMUSIC), and the connection to powerful exact algorithms as subroutines for handable subproblems and other means of hybridization are avenues to be followed.

Applications of metaheuristics are almost uncountable and appear in various journals (e.g., *Journal of Heuristics*), books, and technical reports every day. A helpful source for a subset of successful applications may be special issues of journals or compilations such as [98, 124, 100, 34], just to mention some. Specialized conferences like the *Metaheuristics International Conference* (MIC) are devoted to the topic (see, e.g., [88, 124, 102, 101, 76, 34]) and even more general conferences reveal that metaheuristics have become part of necessary prerequisites for successfully solving optimization problems (see, e.g., [61]). Moreover, ready to use systems such as class libraries and frameworks have been developed, although usually restricted to be applied by the *knowledgeable* user.

Specialized applications also reveal research needs, e.g., in dynamic environments. One example refers to the application of metaheuristics for online optimization; see, e.g., [65].

From a theoretical point of view, the use of most metaheuristics has not yet been fully justified. While convergence results regarding solution quality exist for most metaheuristics once appropriate probabilistic assumptions are made (see, e.g., [66, 8, 42]), these turn out not to be very helpful in practice as usually a disproportionate computation time is required to achieve these results (usually convergence is achieved for the computation time tending to infinity, with a few exceptions, e.g., for the reverse elimination method within tabu search or the pilot method where optimality can be achieved with a finite, but exponential number of steps in the worst case). Furthermore, we have to admit that theoretically one may argue that none of the described metaheuristics is *on average* better than any other. Basically this leaves the choice of a best possible heuristic or related ingredients to the ingenuity of the user/researcher. Some researchers related the acronym of hyper heuristics to the question which (heuristic) method among a given set of methods to choose for a given problem; see, e.g., [20].

Moreover, despite the widespread success of various metaheuristics, researchers occasionally still have a poor understanding of many key theoretical aspects of these algorithms, including models of the high-level run-time dynamics and identification of search space features that influence problem difficulty. Moreover, fitness landscape evaluations are considered in its infancy, too.

From an empirical standpoint it would be most interesting to know which algorithms perform best under various criteria for different classes of problems. Unfortunately, this theme is out of reach as long as we do not have any

well accepted standards regarding the testing and comparison of different methods.

While most papers on metaheuristics claim to provide ‘high quality’ results based on some sort of measure, we still believe that there is a great deal of room for improvement in testing existing as well as new approaches from an empirical point of view (see, e.g., [10, 74, 83]). In a dynamic research process numerical results provide the basis for systematically developing efficient algorithms. The essential conclusions of finished research and development processes should always be substantiated (i.e., empirically and, if necessary, statistically proven) by numerical results based on an appropriate empirical test cycle. Furthermore, even when excellent numerical results are obtained, it may still be possible to compare with a simple random restart procedure and obtain better results in some cases; see, e.g., [62]. However, this comparison is often neglected.

Usually the ways of preparing, performing and presenting experiments and their results are significantly different. The failing of a generally accepted standard for testing and reporting on the testing, or at least a corresponding guideline for designing experiments, unfortunately implies the following observation: Parts of results can be used only in a restricted way, e.g., because relevant data are missing, wrong environmental settings are used, or simply results are glossed over. In the worst case non-sufficiently prepared experiments provide results that are unfit for further use, i.e., any generalized conclusion is out of reach. Future algorithm research needs to provide effective methods for analyzing the performance of, e.g., heuristics in a more scientifically founded way (see, e.g., [127, 4] for some steps into this direction).

A final aspect that deserves special consideration is to investigate the use of information within different metaheuristics. While the adaptive memory programming frame provides a very good entry into this area, this still provides an interesting opportunity to link artificial intelligence with operations research concepts.

References

1. E.H.L. Aarts and J.K. Lenstra, editors. *Local Search in Combinatorial Optimization*. Wiley, Chichester, 1997.
2. E.H.L. Aarts and M. Verhoeven. Local search. In M. Dell’Amico, F. Maffioli, and S. Martello, editors, *Annotated Bibliographies in Combinatorial Optimization*, pages 163–180. Wiley, Chichester, 1997.
3. T. Achterberg and T. Berthold. Improving the feasibility pump. *Discrete Optimization*, 4:77–86, 2007.
4. B. Adenso-Diaz and M. Laguna. Fine-tuning of algorithms using fractional experimental designs and local search. *Operations Research*, 54:99–114, 2006.
5. R.K. Ahuja, O. Ergun, J.B. Orlin, and A.B. Punnen. A survey of very large-scale neighborhood search techniques. *Discrete Applied Mathematics*, 123:75–102, 2002.
6. E. Alba, editor. *Parallel Metaheuristics*. Wiley, Hoboken, 2005.

7. E. Alba and R. Marti, editors. *Metaheuristic Procedures for Training Neural Networks*. Springer, New York, 2006.
8. I. Althöfer and K.-U. Koschnick. On the convergence of ‘threshold accepting’. *Applied Mathematics and Optimization*, 24:183–195, 1991.
9. T. Bäck, D.B. Fogel, and Z. Michalewicz, editors. *Handbook of Evolutionary Computation*. Institute of Physics Publishing, Bristol, 1997.
10. R.S. Barr, B.L. Golden, J.P. Kelly, M.G.C. Resende, and W.R. Stewart. Designing and reporting on computational experiments with heuristic methods. *Journal of Heuristics*, 1:9–32, 1995.
11. M.B. Bastos and C.C. Ribeiro. Reactive tabu search with path relinking for the Steiner problem in graphs. In C.C. Ribeiro and P. Hansen, editors, *Essays and Surveys in Metaheuristics*, pages 39–58. Kluwer, Boston, 2002.
12. R. Battiti. Machine learning methods for parameter tuning in heuristics. Position paper for the 5th DIMACS Challenge Workshop: Experimental Methodology Day, 1996.
13. R. Battiti. Reactive search: Toward self-tuning heuristics. In V.J. Rayward-Smith, I.H. Osman, C.R. Reeves, and G.D. Smith, editors, *Modern Heuristic Search Methods*, pages 61–83. Wiley, Chichester, 1996.
14. R. Battiti, M. Brunato, and F. Mascia. *Reactive Search and Intelligent Optimization*. Springer, New York, 2009.
15. R. Battiti and G. Tecchiolli. The reactive tabu search. *ORSA Journal on Computing*, pages 126–140, 1994.
16. L. Bertacco, M. Fischetti, and A. Lodi. A feasibility pump heuristic for general mixed integer problems. *Discrete Optimization*, 4(1):77–86, 2007.
17. D.P. Bertsekas, J.N. Tsitsiklis, and C. Wu. Rollout algorithms for combinatorial optimization. *Journal of Heuristics*, 3:245–262, 1997.
18. C. Bierwirth, D.C. Mattfeld, and J.P. Watson. Landscape regularity and random walks for the job-shop scheduling problem. In J. Gottlieb and G.R. Raidl, editors, *Evolutionary Computation in Combinatorial Optimization, 4th European Conference, EvoCOP 2004*, volume 3004 of *Lecture Notes in Computer Science*, pages 21–30. Springer, 2004.
19. C. Blum and A. Roli. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Computing Surveys*, 35:268–308, 2003.
20. E.K. Burke, G. Kendall, J. Newall, E. Hart, P. Ross, and S. Schulenburg. Hyperheuristics: An emerging direction in modern search technology. In F.W. Glover and G.A. Kochenberger, editors, *Handbook of Metaheuristics*, pages 457–474. Kluwer, Boston, 2003.
21. M. Caserta and E. Quiñonez Rico. A cross entropy-lagrangean hybrid algorithm for the multi-item capacitated lot sizing problem with setup times. *Computers & Operations Research*, 36(2):530–548, 2009.
22. R. Cerulli, A. Fink, M. Gentili, and S. Voß. Extensions of the minimum labelling spanning tree problem. *Journal of Telecommunications and Information Technology*, 4/2006:39–45, 2006.
23. I. Charon and O. Hudry. The noising method: A new method for combinatorial optimization. *Operations Research Letters*, 14:133–137, 1993.
24. C. Cotta and A. Fernández. Analyzing fitness landscapes for the optimal golomb ruler problem. In G.R. Raidl and J. Gottlieb, editors, *Evolutionary Computation in Combinatorial Optimization, 5th European Conference, EvoCOP 2005*, volume 3448 of *Lecture Notes in Computer Science*, pages 68–79. Springer, 2005.
25. S. P. Coy, B.L. Golden, G.C. Rungen, and E.A. Wasil. Using experimental design to find effective parameter settings for heuristics. *Journal of Heuristics*, 7:77–97, 2000.
26. T.G. Crainic, M. Toulouse, and M. Gendreau. Toward a taxonomy of parallel tabu search heuristics. *INFORMS Journal on Computing*, 9:61–72, 1997.
27. E. Danna, E. Rothberg, and C. Le Pape. Exploring relaxation induced neighborhoods to improve MIP solutions. *Mathematical Programming A*, 102:71–90, 2005.

28. P. De Boer, D.P. Kroese, S. Mannor, and R.Y. Rubinstein. A tutorial on the cross-entropy method. *Annals of Operations Research*, 134:19–67, 2005.
29. J. Demšar. Statistical comparison of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7:1–30, 2006.
30. L. Di Gaspero and A. Schaerf. EASYLOCAL++: An object-oriented framework for the flexible design of local-search algorithms. *Software – Practice and Experience*, 33:733–765, 2003.
31. T.G. Dietterich. Approximate statistical test for comparing supervised classification learning algorithms. *Neural Computation*, 10(7):1895–1923, 1998.
32. M. Dorigo, V. Maniezzo, and A. Colorni. Ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man and Cybernetics*, B - 26:29–41, 1996.
33. M. Dorigo and T. Stützle. *Ant Colony Optimization*. MIT Press, Cambridge, 2004.
34. K.F. Dörner, M. Gendreau, P. Greistorfer, W.J. Gutjahr, R.F. Hartl, and M. Reimann, editors. *Metaheuristics: Progress in Complex Systems Optimization*. Springer, New York, 2007.
35. K.A. Dowsland. Simulated annealing. In C. Reeves, editor, *Modern Heuristic Techniques for Combinatorial Problems*, pages 20–69. Halsted, Blackwell, 1993.
36. J. Dreo, A. Petrowski, P. Siarry, and E. Taillard. *Metaheuristics for Hard Optimization*. Springer, Berlin, 2006.
37. G. Dueck and T. Scheuer. Threshold accepting: a general purpose optimization algorithm appearing superior to simulated annealing. *Journal of Computational Physics*, 90:161–175, 1990.
38. C.W. Duin and S. Voß. Steiner tree heuristics - a survey. In H. Dyckhoff, U. Derigs, M. Salomon, and H.C. Tijms, editors, *Operations Research Proceedings 1993*, pages 485–496, Berlin, 1994. Springer.
39. C.W. Duin and S. Voß. The pilot method: A strategy for heuristic repetition with application to the Steiner problem in graphs. *Networks*, 34:181–191, 1999.
40. J. Eckstein and M. Nediak. Pivot, cut, and dive: a heuristic for 0-1 mixed integer programming. *Journal of Heuristics*, 13:471–503, 2007.
41. M. Ehrgott and X. Gandibleux. Bound sets for biobjective combinatorial optimization problems. *Computers & Operations Research*, 34(9):2674–2694, 2007.
42. U. Faigle and W. Kern. Some convergence results for probabilistic tabu search. *ORSA Journal on Computing*, 4:32–37, 1992.
43. P. Festa and M.G.C. Resende. An annotated bibliography of GRASP. Technical report, AT&T Labs Research, 2004.
44. G.R. Filho and L.A. Lorena. Constructive genetic algorithm and column generation: an application to graph coloring. In *Proceedings of APORS 2000 - The Fifth Conference of the Association of Asian-Pacific Operations Research Society within IFORS 2000*.
45. A. Fink and S. Voß. HOTFRAME: A heuristic optimization framework. In S. Voß and D.L. Woodruff, editors, *Optimization Software Class Libraries*, pages 81–154. Kluwer, Boston, 2002.
46. M. Fischetti, F. Glover, and A. Lodi. The feasibility pump. *Mathematical Programming*, A 104:91–104, 2005.
47. M. Fischetti and A. Lodi. Local branching. *Mathematical Programming*, B 98:23–47, 2003.
48. D.B. Fogel. On the philosophical differences between evolutionary algorithms and genetic algorithms. In D.B. Fogel and W. Atmar, editors, *Proceedings of the Second Annual Conference on Evolutionary Programming*, pages 23–29. Evolutionary Programming Society, La Jolla, 1993.
49. D.B. Fogel. *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*. IEEE Press, New York, 1995.
50. A. M. Glenberg. *Learning from Data: An Introduction to Statistical Reasoning*. Lawrence Erlbaum Associates, Mahwah, New Jersey, 1996.

51. F. Glover. Heuristics for integer programming using surrogate constraints. *Decision Sciences*, 8:156–166, 1977.
52. F. Glover. Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, 13:533–549, 1986.
53. F. Glover. Tabu search – Part II. *ORSA Journal on Computing*, 2:4–32, 1990.
54. F. Glover. Scatter search and star-paths: beyond the genetic metaphor. *OR Spektrum*, 17:125–137, 1995.
55. F. Glover. Tabu search and adaptive memory programming – Advances, applications and challenges. In R.S. Barr, R.V. Helgason, and J.L. Kennington, editors, *Advances in Metaheuristics, Optimization, and Stochastic Modeling Technologies*, pages 1–75. Kluwer, Boston, 1997.
56. F. Glover and M. Laguna. General purpose heuristics for integer programming - part I. *Journal of Heuristics*, 2(4):343–358, 1997.
57. F. Glover and M. Laguna. General purpose heuristics for integer programming - part II. *Journal of Heuristics*, 3(2):161–179, 1997.
58. F. Glover and M. Laguna. *Tabu Search*. Kluwer, Boston, 1997.
59. F.W. Glover and G.A. Kochenberger, editors. *Handbook of Metaheuristics*. Kluwer, Boston, 2003.
60. D.E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, 1989.
61. B.L. Golden, S. Raghavan, and E.A. Wasil, editors. *The Next Wave in Computing, Optimization, and Decision Technologies*. Kluwer, Boston, 2005.
62. A.M. Gomes and J.F. Oliveira. Solving irregular strip packing problems by hybridising simulated annealing and linear programming. *European Journal of Operational Research*, 171:811–829, 2006.
63. P. Greistorfer, A. Lokketangen, D.L. Woodruff, and S. Voß. Sequential versus simultaneous maximization of objective and diversity. *Journal of Heuristics*, 14:613–625, 2008.
64. P. Greistorfer and S. Voß. Controlled pool maintenance for meta-heuristics. In C. Rego and B. Alidaee, editors, *Metaheuristic Optimization via Memory and Evolution*, pages 387–424. 2005.
65. K. Gutenschwager, C. Niklaus, and S. Voß. Dispatching of an electric monorail system: Applying meta-heuristics to an online pickup and delivery problem. *Transportation Science*, 38:434–446, 2004.
66. B. Hajek. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, 13:311–329, 1988.
67. P. Hansen, V. Maniezzo, and S. Voß. Special issue on mathematical contributions to metaheuristics editorial. *Journal of Heuristics*, 15(3):197–199, 2009.
68. P. Hansen and N. Mladenović. An introduction to variable neighborhood search. In S. Voß, S. Martello, I.H. Osman, and C. Roucairol, editors, *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*, pages 433–458. Kluwer, Boston, 1999.
69. P. Hansen, N. Mladenović, and D. Perez-Brito. Variable neighborhood decomposition search. *Journal of Heuristics*, 7(4):335–350, 2001.
70. J.P. Hart and A.W. Shogan. Semi-greedy heuristics: An empirical study. *Operations Research Letters*, 6:107–114, 1987.
71. A. Hertz and D. Kobler. A framework for the description of evolutionary algorithms. *European Journal of Operational Research*, 126:1–12, 2000.
72. F. Hoffmeister and T. Bäck. Genetic algorithms and evolution strategies: Similarities and differences. In H.-P. Schwefel and R. Männer, editors, *Parallel Problem Solving from Nature*, volume 496 of *Lecture Notes in Computer Science*, pages 455–469. Springer, 1991.
73. J.H. Holland. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor, 1975.

74. J.N. Hooker. Testing heuristics: We have it all wrong. *Journal of Heuristics*, 1:33–42, 1995.
75. H.H. Hoos and T. Stützle. *Stochastic Local Search – Foundations and Applications*. Elsevier, Amsterdam, 2005.
76. T. Ibaraki, K. Nonobe, and M. Yagiura, editors. *Metaheuristics: Progress as Real Problem Solvers*. Springer, New York, 2005.
77. L. Ingber. Adaptive simulated annealing (ASA): Lessons learned. *Control and Cybernetics*, 25:33–54, 1996.
78. A. Jaszkiwicz. A comparative study of multiple-objective metaheuristics on the bi-objective set covering problem and the pareto memetic algorithm. *Annals of Operations Research*, 131:215–235, 2004.
79. D.S. Johnson, C.R. Aragon, L.A. McGeoch, and C. Schevon. Optimization by simulated annealing: An experimental evaluation; part i, graph partitioning. *Operations Research*, 37:865–892, 1989.
80. S. Kirkpatrick, C.D. Gelatt Jr., and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.
81. M. Laguna and R. Martí. *Scatter Search*. Kluwer, Boston, 2003.
82. S. Lin and B.W. Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations Research*, 21:498–516, 1973.
83. C. McGeoch. Toward an experimental method for algorithm simulation. *INFORMS Journal on Computing*, 8:1–15, 1996.
84. C. Meloni, D. Pacciarelli, and M. Pranzo. A rollout metaheuristic for job shop scheduling problems. *Annals of Operations Research*, 131:215–235, 2004.
85. Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer, Berlin, 3 edition, 1999.
86. Z. Michalewicz and D.B. Fogel. *How to Solve It: Modern Heuristics*. Springer, Berlin, 2 edition, 2004.
87. P. Moscato. An introduction to population approaches for optimization and hierarchical objective functions: A discussion on the role of tabu search. *Annals of Operations Research*, 41:85–121, 1993.
88. I.H. Osman and J.P. Kelly, editors. *Meta-Heuristics: Theory and Applications*. Kluwer, Boston, 1996.
89. M.W. Park and Y.D. Kim. A systematic procedure for setting parameters in simulated annealing algorithms. *Computers & Operations Research*, 25(3):207–217, 1998.
90. R. Parson and M.E. Johnson. A case study in experimental design applied to genetic algorithms with applications to DNA sequence assembly. *American Journal of Mathematical and Management Sciences*, 17(3):369–396, 1997.
91. J. Pearl. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley, Reading, 1984.
92. E. Pesch and F. Glover. TSP ejection chains. *Discrete Applied Mathematics*, 76:165–182, 1997.
93. G. Polya. *How to solve it*. Princeton University Press, Princeton, 1945.
94. J. Puchinger and G.R. Raidl. An evolutionary algorithm for column generation in integer programming: an effective approach for 2D bin packing. In X. Yao, E.K. Burke, J.A. Lozano, J. Smith, J.J. Merelo-Guervos, J.A. Bullinaria, J.E. Rowe, P. Tino, A. Kaban, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature – PPSN VIII*, volume 3242 of *Lecture Notes in Computer Science*, pages 642–651. Springer Verlag, 2004.
95. J. Puchinger and G.R. Raidl. Combining metaheuristics and exact algorithms in combinatorial optimization: A survey and classification. In J. Mira and J.R. Álvarez, editors, *Proceedings of the First International Work-Conference on the Interplay Between Natural and Artificial Computation, Part II*, volume 3562 of *Lecture Notes in Computer Science*, pages 41–53. Springer, 2005.

96. G.R. Raidl. A unified view on hybrid metaheuristics. In F. Almeida, M.J. Blesa, C. Blum, J.M. Moreno-Vega, M.M. Pérez, A. Roli, and M. Sampels, editors, *Hybrid Metaheuristics*, volume 4030 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2006.
97. B. Rangaswamy, A. S. Jain, and F. Glover. Tabu search candidate list strategies in scheduling. pages 215–233, 1998.
98. V.J. Rayward-Smith, I.H. Osman, C.R. Reeves, and G.D. Smith, editors. *Modern Heuristic Search Methods*. Wiley, Chichester, 1996.
99. C.R. Reeves and J.E. Rowe. *Genetic Algorithms: Principles and Perspectives*. Kluwer, Boston, 2002.
100. C. Rego and B. Alidaee, editors. *Metaheuristic Optimization via Memory and Evolution*. 2005.
101. M.G.C. Resende and J.P. de Sousa, editors. *Metaheuristics: Computer Decision-Making*. Kluwer, Boston, 2004.
102. C.C. Ribeiro and P. Hansen, editors. *Essays and Surveys in Metaheuristics*. Kluwer, Boston, 2002.
103. F. Rossi, P. van Beek, and T. Walsh, editors. *Handbook of Constraint Programming (Foundations of Artificial Intelligence)*. Elsevier, 2006.
104. R.Y. Rubinstein. Optimization of Computer Simulation Models with Rare Events. *European Journal of Operational Research*, 99:89–112, 1997.
105. M. Sakawa. *Genetic Algorithms and Fuzzy Multiobjective Optimization*. Kluwer, Boston, 2001.
106. H.-P. Schwefel and T. Bäck. Artificial evolution: How and why? In D. Quagliarella, J. Périaux, C. Poloni, and G. Winter, editors, *Genetic Algorithms and Evolution Strategy in Engineering and Computer Science: Recent Advances and Industrial Applications*, pages 1–19. Wiley, Chichester, 1998.
107. P. Shaw. Using constraint programming and local search methods to solve vehicle routing problems. Working paper, ILOG S.A., Gently, France, 1998.
108. K. Smith. Neural networks for combinatorial optimisation: A review of more than a decade of research. *INFORMS Journal on Computing*, 11:15–34, 1999.
109. M. Sniedovich and S. Voß. The corridor method: A dynamic programming inspired metaheuristic. *Control and Cybernetics*, 35:551–578, 2006.
110. L. Sondergeld. *Performance Analysis Methods for Heuristic Search Optimization with an Application to Cooperative Agent Algorithms*. Shaker, Aachen, 2001.
111. R.H. Storer, S.D. Wu, and R. Vaccari. Problem and heuristic space search strategies for job shop scheduling. *ORSA Journal on Computing*, 7:453–467, 1995.
112. E. Taillard and S. Voß. POPMUSIC — partial optimization metaheuristic under special intensification conditions. In C.C. Ribeiro and P. Hansen, editors, *Essays and Surveys in Metaheuristics*, pages 613–629. Kluwer, Boston, 2002.
113. E. Taillard, P. Waelti, and J. Zuber. Few statistical tests for proportions comparison. *European Journal of Operational Research*, 185(3):1336–1350, 2006.
114. É.D. Taillard, L.M. Gambardella, M. Gendreau, and J.-Y. Potvin. Adaptive memory programming: A unified view of meta-heuristics. *European Journal of Operational Research*, 135:1–16, 2001.
115. J. Tavares, F. Pereira, and E. Costa. Multidimensional knapsack problem: A fitness landscape analysis. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 38(3):604–616, 2008.
116. R.J.M. Vaessens, E.H.L. Aarts, and J.K. Lenstra. A local search template. *Computers & Operations Research*, 25:969–979, 1998.
117. M.G.A. Verhoeven and E.H.L. Aarts. Parallel local search techniques. *Journal of Heuristics*, 1:43–65, 1995.
118. S. Voß. *Intelligent Search*. Manuscript, TU Darmstadt, 1993.
119. S. Voß. Tabu search: applications and prospects. In D.-Z. Du and P. Pardalos, editors, *Network Optimization Problems*, pages 333–353. World Scientific, Singapore, 1993.

120. S. Voß. Observing logical interdependencies in tabu search: Methods and results. In V.J. Rayward-Smith, I.H. Osman, C.R. Reeves, and G.D. Smith, editors, *Modern Heuristic Search Methods*, pages 41–59, Chichester, 1996. Wiley.
121. S. Voß. Meta-heuristics: The state of the art. In A. Nareyek, editor, *Local Search for Planning and Scheduling*, volume 2148 of *Lecture Notes in Artificial Intelligence*, pages 1–23. Springer, 2001.
122. S. Voß. Metaheuristics. In C.A. Floudas and P.M. Pardalos, editors, *Encyclopedia of Optimization*. Springer, New York, 2008.
123. S. Voß, A. Fink, and C. Duin. Looking ahead with the pilot method. *Annals of Operations Research*, 136:285–302, 2005.
124. S. Voß, S. Martello, I.H. Osman, and C. Roucairol, editors. *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*. Kluwer, Boston, 1999.
125. S. Voß and D.L. Woodruff, editors. *Optimization Software Class Libraries*. Kluwer, Boston, 2002.
126. J. P. Watson, L. D. Whitley, and A. E. Howe. Linking search space structure, run-time dynamics, and problem difficulty: A step toward demystifying tabu search. *Journal of Artificial Intelligence Research*, 24:221–261, 2005.
127. D. Whitley, S. Rana, J. Dzuberka, and K.E. Mathias. Evaluating evolutionary algorithms. *Artificial Intelligence*, 85:245–276, 1996.
128. D.H. Wolpert and W.G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1:67–82, 1997.
129. D.L. Woodruff. Proposals for chunking and tabu search. *European Journal of Operational Research*, 106:585–598, 1998.
130. D.L. Woodruff. A chunking based selection strategy for integrating meta-heuristics with branch and bound. In S. Voß, S. Martello, I.H. Osman, and C. Roucairol, editors, *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*, pages 499–511. Kluwer, Boston, 1999.
131. J. Xu, S.Y. Chiu, and F. Glover. Fine-tuning a tabu search algorithm with statistical tests. *International Transactions in Operational Research*, 5(3):233–244, 1998.
132. J.H. Zar. *Biostatistical Analysis*. Prentice Hall, Upper Saddle River, New Jersey, 1999.
133. M. Zlochin, M. Birattari, N. Meuleau, and M. Dorigo. Model-based search for combinatorial optimization. *Annals of Operations Research*, 131(1):373–395, 2004.