
Ciplus
Band 4/2018

A new Taxonomy of Continuous Global Optimization Algorithms

Jörg Stork, A.E. Eiben, and Thomas Bartz-Beielstein

Technology
Arts Sciences
TH Köln

A new Taxonomy of Continuous Global Optimization Algorithms

Jörg Stork¹, A.E. Eiben², and Thomas Bartz-Beielstein¹

¹TH Köln, Steinmüellerallee 1, 51648 Gummersbach, Germany

²VU Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, Netherlands

joerg.stork@th-koeln.de, a.e.eiben@vu.nl, thomas.bartz-beielstein@th-koeln.de

Abstract

Surrogate-based optimization and nature-inspired metaheuristics have become the state of the art in solving real-world optimization problems. Still, it is difficult for beginners and even experts to get an overview that explains their advantages in comparison to the large number of available methods in the scope of continuous optimization. Available taxonomies lack the integration of surrogate-based approaches and thus their embedding in the larger context of this broad field. This article presents a taxonomy of the field, which further matches the idea of nature-inspired algorithms, as it is based on the human behavior in path finding. Intuitive analogies make it easy to conceive the most basic principles of the search algorithms, even for beginners and non-experts in this area of research. However, this scheme does not oversimplify the high complexity of the different algorithms, as the class identifier only defines a descriptive meta-level of the algorithm search strategies. The taxonomy was established by exploring and matching algorithm schemes, extracting similarities and differences, and creating a set of classification indicators to distinguish between five distinct classes. In practice, this taxonomy allows recommendations for the applicability of the corresponding algorithms and helps developers trying to create or improve their own algorithms.

1 Introduction

Continuous global optimization (CGO) tackles various difficult problems emerging from the context of complex physical or chemical processes. Solving optimization problems of this kind necessarily relies on performing real-world experiments or on using computer simulations, which are frequently employed in black-box fashions. A fundamental challenge in such systems is the high cost of function evaluations. Whether we are probing the real physical system or querying the simulator, the time needed to receive an objective function value is typically very high and can range from hours to months. CGO methods for such problems thus need to fulfill a certain set of requirements. They need to work with black-box style probes only, so without any further information on the structure of the problem. Further they must approach the vicinity of the global optimum with a limited number of function evaluations.

The improvement of computational power in the last decades has been influencing the development of algorithms. A massive amount of computational power became available for

researchers worldwide through multi-core desktop machines, parallel computing, and high-performance computing clusters. This has improved the following fields of research: Firstly, the development of more complex, nature-inspired, and generally applicable heuristics, so called metaheuristics. Secondly, major advances in the field of accurate, data-driven approximation models, so-called surrogate models and their embodiment in an optimization process.

Nowadays, CGO differs largely from early approaches. For example, multi-staged methods evaluate objective functions not directly on the problem. They utilize a combination of surrogate modeling with classical or metaheuristic optimization methods to maximize the use of available problem information instead. These frameworks, such as *sequential parameter optimization* (Bartz-Beielstein, Lasarczyk, and Preuß, 2005) or the *surrogate management framework* (Booker, Dennis Jr, Frank, Serafini, Torczon, and Trosset, 1999; Serafini, 1999), define a new class of algorithms that are not well integrated in previous taxonomies.

In this article we will propose a new taxonomy on basis of algorithm features and display plausible descriptions founded on the natural human behavior in path finding. To establish a comprehensive taxonomy, we focus on identifying key elements of algorithm design and utilizing these to define a clear separation between a small number of algorithm classes. Although abstraction is necessary for developing our results, we will present results that are useful for practitioners. The utilized abstraction allows us to present simply comprehensible ideas on how the individual classes differ and moreover, how the respective algorithms perform their search. For this purpose, we divide CGO algorithms into five intuitive classes: *Exact*, *Wanderer*, *Guide*, *Cartographer*, and *Hybrid*. This article particularly addresses different kinds of readers: beginners will find an intuitive taxonomy of CGO algorithms, especially with regard to common metaheuristics and newer developments in the field of surrogate-based optimization. For advanced readers, we also discuss the suitability of certain classes for specific problem properties to provide basic knowledge for reasonable algorithm selection. An extensive list of references is provided for the experienced users. The taxonomy can be used to create realistic comparisons and benchmarks for the different classes of algorithms. It further provides insights for users, who aim to develop new search strategies, operators and algorithms.

The goal of global optimization is to find the overall best solution, i.e., for the common task of minimization, to discover decision variable values which minimize the objective function value. We denote the global search space as compact set $\mathcal{S} = \{\mathbf{x} \mid \mathbf{x}_l \leq \mathbf{x} \leq \mathbf{x}_u\}$ with $\mathbf{x}_l, \mathbf{x}_u \in \mathbb{R}^n$ being the explicit, finite lower and upper bounds on \mathbf{x} .

Given a real valued objective function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ with real valued input vectors \mathbf{x} we attempt to find the location $\mathbf{x} \in \mathbb{R}^n$ which minimizes the function: $\arg \min f(\mathbf{x}), \mathbf{x} \in \mathcal{S}$.

Finding the global optimum is always the ultimate goal and as such desirable, but for many practical problems a solution improving the current best solution in a given budget of evaluations or time will still be a success. Particular in CGO the global optimum commonly cannot be identified exactly, thus modern heuristics are designed to spend their resources as efficiently as possible to approximate near-best solutions, while finding the global optimum is never guaranteed. The remainder of this article is organized as follows: Section 2 presents the development of optimization algorithms and their core concepts. Section 3 motivates a new taxonomy by reviewing the history of available CGO taxonomies, illustrates algorithm design aspects and presents extracted classification features. Sections 4 to 8 introduce the five different classes of the new taxonomy with examples and suggestions regarding their applicability. Section 9 summarizes and concludes the article with the recent trends and challenges in CGO and currently important research fields.

2 Evolution of Optimization Algorithms

In order to develop a taxonomy, it is necessary to understand the methodology and development history of the corresponding algorithms. Before presenting the requirements for the new taxonomy in Section 3, we will describe the fundamental principles of modern search algorithms, particular the elements and backgrounds of surrogate-based optimization.

2.1 Heuristics and Metaheuristics

In modern computer-aided optimization, heuristics and metaheuristics are well established solution techniques. Although presenting solutions which are not guaranteed to be optimal or perfect, their general applicability and ability of presenting fast sufficient solutions make them very attractive for applied optimization, particular for industrial problems. They are built upon the principle of *trial and error*, where solution candidates are evaluated and rewarded with a *fitness*. The term *fitness* has its origins in evolutionary computation (Eiben and Smith, 2015), where the fitness describes the competitive ability of an individual in the reproduction process. The fitness is in its simplest form the objective function value $y = f(\mathbf{x})$ in relation to the optimization goal, e.g., in a minimization problem smaller values have a higher fitness. Moreover, it can be part of the search strategy, e.g., scaled or adjusted by additional functions, particular for multi-objective or constrained optimization.

Heuristics can be defined as problem-dependent algorithms, which are developed or adapted to the particularities of a specific optimization problem or problem instance (Pearl, 1985). Typically, heuristics perform evaluations in a systematic manner, although utilizing stochastic elements. Heuristics use this principle to provide fast, not necessarily exact (i.e., not optimal) numerical solutions to optimization problems. Moreover, heuristics are often greedy to provide fast solutions, but get trapped in local optima and fail to find the global optimum.

In their starting days in the 1960s, heuristics were not considered as reliable problem solvers, because most researchers in academia preferred classical mathematical approaches and only some practitioners used heuristics to get fast, possibly inaccurate solutions (Zanakis and Evans, 1981). This situation changed in the 1970s. Heuristic optimization became a major part of academic research. Possible reasons for this change might be:

- The need to solve more sophisticated nondeterministic problems, which had a polynomial runtime and thus could not be solved efficiently with exact algorithms (Fomin and Kaski, 2013).
- The availability and easy access for academics to more computational power.

Further advantages of heuristics were summarized by Zanakis and Evans (1981). The most important are:

- Simplicity of the algorithm.
- Accuracy, i.e., small error of final solution.
- Robustness, i.e., good solutions within reasonable time for different problems.
- Speed, i.e., duration of the computation.

While heuristics are developed and optimized to efficiently solve a certain problem, the improved availability of computer resources gave rise to higher-level heuristics, the metaheuristics. *Metaheuristics* can be defined as problem independent, general purpose optimization algorithms. They are applicable to a wide range of problems and problem instances. The term *meta* describes the higher-level general methodology, which is utilized to guide the underlying heuristic strategy (Talbi, 2009).

They share the following characteristics (Boussaïd, Lepagnot, and Siarry, 2013):

- The algorithms are nature-inspired; they follow certain principles from natural phenomena or behaviors (e.g., biological evolution, physics, social behavior).
- The search process involves stochastic parts; it is based on probability distributions and random processes.
- They do not use the gradient or Hessian of the objective function or rely on information of the process which is available before the start of the optimization run, so-called *a priori* information.
- As they are meant to be general applicable solvers, they include a set of control parameters to adjust the search strategy.

During the remainder of this article we will focus on heuristic, respectively metaheuristic algorithms.

2.2 Modern Optimization Algorithms

Based on the fundamentals of heuristics and metaheuristics, we are able to identify a similarity in the design of modern optimization algorithms, which targets a large class of problems. Importantly, we have to consider the *No Free Lunch Theorem* (Wolpert and Macready, 1997), which states that there is no optimization algorithm that is superior to all others if their performance is averaged over all possible problems. Consequently, any algorithm needs to be adapted to the structure of the problem at hand to achieve optimal performance. This can be considered during the construction of an algorithm, before the optimization by parameter *tuning* or during the run by parameter *control* (Bartz-Beielstein et al., 2005; Eiben, Hinterding, and Michalewicz, 1999). Törn and Zilinskas (1989) mention three principles for the construction of an optimization algorithm:

1. An algorithm utilizing all available *a priori* information will outperform a method using less information.
2. If no *a priori* information is available, the information is completely based on evaluated candidate points and their objective values.
3. Given a fixed number of evaluated points, optimization algorithms will only differ from each other in the distribution of candidate points.

As most modern algorithms focus on handling problems where little or no *a priori* information is given, the principles displayed above lead to the conclusion that the most crucial design aspect of any algorithm is to find a strategy to distribute the initial candidates in the search space and to generate new candidates based on variation of solutions. These two procedures define the *search strategy*, which needs to follow the two competing goals of *exploration* and *exploitation*. In general, the main goal of any method is to reach their target with high efficiency, i.e., to discover optima fast and accurate with as little resources as possible. Moreover, the goal is not mandatory finding the global optimum, which is a demanding and expensive task for many problems, but to identify a valuable local optimum or to improve the current available solution. We will explicitly discuss the design of modern optimization algorithms in Section 3.3.

2.3 Exact Algorithms

Exact algorithms, also referred to as *complete* algorithms (Neumaier, 2004), are a special class of *deterministic, systematic* and *non-heuristic* optimization algorithms. If sufficient *a priori* information about the objective function is available, they have a guarantee of finding the global optimum within using a predictable amount of resources, such as function evaluations or computation time (Fomin and Kaski, 2013). If they are applicable to the problem, these algorithms are more reliable than heuristics, as they allow convergence proofs of finding the global optimum. Without available *a priori* information, the stopping criterium needs to be defined by a heuristic approach, which softens the guarantee of finding the optimum. Moreover, it is theoretical possible to apply these algorithms to the class of black-box problems with the given ability to find the global optimum with certainty after finite time. However, they will need exponential computation time due to an expensive, dense search. This renders them not applicable to many resource-limited applications. The *exact class*, presented in Section 4, contains the related algorithms.

2.4 Surrogate-based Optimization Algorithms

Surrogate-based optimization algorithms are designed to process expensive and complex problems, which arise from real-world applications and sophisticated computational models. Real-world problems are commonly black-box, which means that they only provide very sparse domain knowledge. Consequently, problem information needs to be exploited by experiments or function evaluations.

Surrogate-based optimization is developed to optimally exhaust the available information by utilizing a *surrogate model*. A surrogate model is an approximation which substitutes the original expensive objective function, real-world process, physical simulation, or computational process during the optimization. In general, surrogates are either simplified *physical* or numerical models based on knowledge about the physical system, or empirical *functional* models based on knowledge acquired from evaluations and sparse sampling of the parameter space (Søndergaard, Madsen, and Nielsen, 2003). In this work, we focus on the latter. The terms *surrogate model*, *meta-model*, *response surface model* and *posterior distribution* are used synonymous in the common literature (Mockus, 1974; Jones, 2001; Bartz-Beielstein and Zaefferer, 2017). We will briefly refer to a surrogate model as *surrogate*. Furthermore, we assume that it is crucial to distinguish between the use of an explicit surrogate of the objective function and general *model-based* optimization (Zlochín, Birattari, Meuleau, and Dorigo, 2004), which additionally refers to methods, where a statistical model is used to generate new candidate solutions (cf. Section 3.3). As these two definitions of model-based optimization are frequently used in a non-consistent manner, we will clearly distinguish between the two different terms *surrogate-based* and *model-based* to avoid confusions. Another term present in the literature is *surrogate-assisted* optimization, which mostly refers to the application of surrogates in population-based evolutionary computation (Jin, 2011), where evolutionary optimization and surrogate-based optimization are applied in an hybrid approach (see Section 8).

Important publications featuring overviews or surveys on surrogates and surrogate-based optimization were presented by Sacks, Welch, Mitchell, and Wynn (1989), Jones (2001), Queipo, Haftka, Shyy, Goel, Vaidyanathan, and Tucker (2005), Forrester and Keane (2009). Surrogate-based optimization is commonly applied, but not limited to the case of complex real-world optimization applications, where two typical problem layers and a surrogate layer can be de-

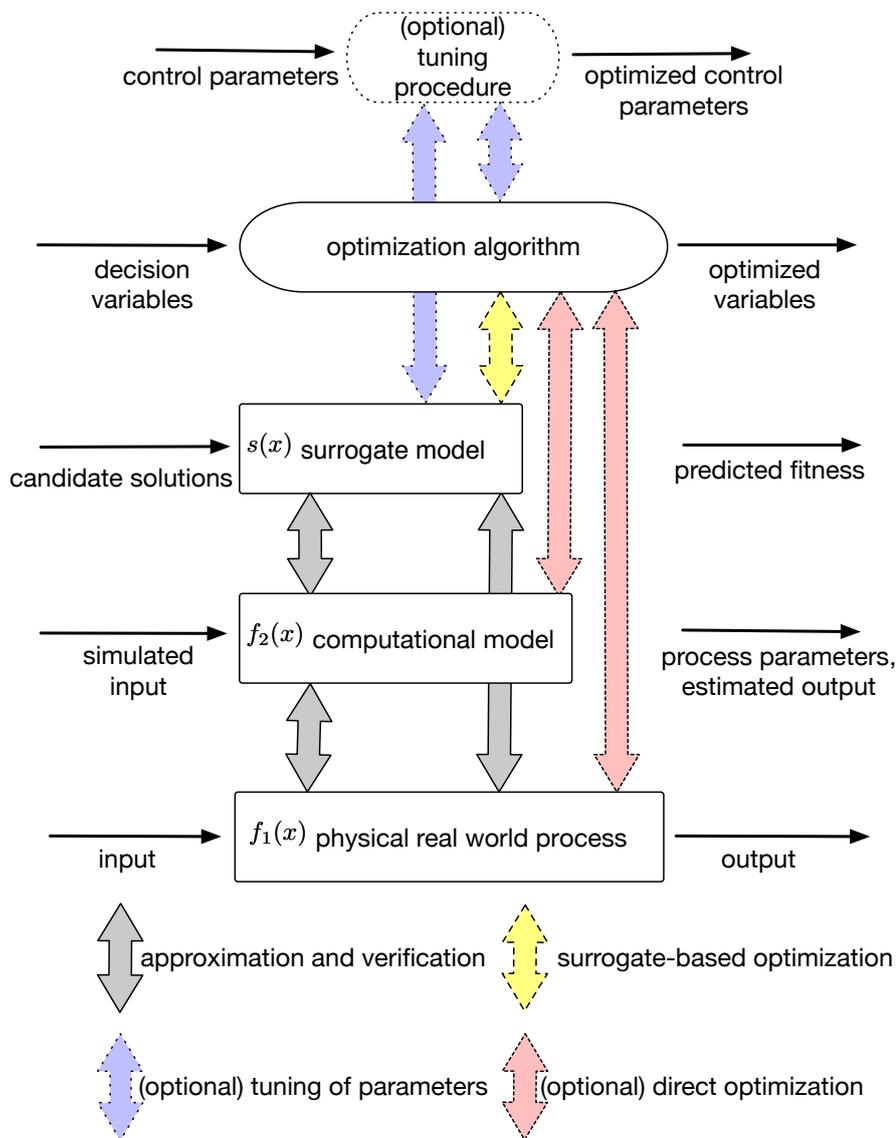


Figure 1: A surrogate-based optimization process of a real-world process with the different objective function layers and outlined inputs and outputs; the complexity of the objective functions is visible by the decreasing size of their boxes. The full grey arrows illustrate the approximation and verification paths, the yellow dashed and red dotted arrows indicate the surrogate-based optimization and optional direct optimization. The blue dashed arrows show the optional parameter tuning of the optimization algorithm or surrogate modeling process.

defined. The defined layers can be transferred to different computational problems with expensive function evaluations, such as complex algorithms or machine learning tasks. Each layer can be the target of an optimization or used to retrieve information to guide the optimization process. Figure 1 illustrates the different layers of objective functions and the surrogate-based optimization process for real-world problems. In this case, the objective function layers, from the bottom up, are:

- L1 The real-world application $f_1(\mathbf{x})$, given by the physical process itself or a physical model. Direct optimization is often expensive or even impossible, due to evaluations involving resource demanding prototype building or even hazardous experiments.
- L2 The computational model $f_2(\mathbf{x})$, given by a simulation of the physical process or a complex computational model, e.g., a computational fluid dynamics model or structural dynamics model. A single computation may take minutes, hours, or even weeks to compute.
- L3 The surrogate $s(\mathbf{x})$, given by a data-driven regression model. The accuracy heavily depends on the underlying surrogate type and number of available information (i.e., function evaluations). The optimization is, compared to the other layers, typically cheap. Surrogates are constructed either for the real-world application $f_1(\mathbf{x})$ or the computational model $f_2(\mathbf{x})$.

Furthermore, the surrogate-based optimization cycle includes the optimization process itself, which is given by an adequate optimization algorithm for the selected objective function layer. No surrogate-based optimization is performed, if the optimization is directly applied to $f_1(\mathbf{x})$ or $f_2(\mathbf{x})$. The surrogate-based optimization uses $f_1(\mathbf{x})$ or $f_2(\mathbf{x})$ for verification of promising solution candidates. Moreover, the control parameters of the optimization algorithm or even the complete optimization cycle including the surrogate modeling process can be tuned (Bartz-Beielstein et al., 2005).

Each layer imposes different evaluation costs and solution accuracies: the real-world problem is the most expensive to evaluate, while the surrogate is the cheapest to evaluate. The main benefit of using surrogates is thus the reduction of needed expensive function evaluations on the objective function $f_1(\mathbf{x})$ or $f_2(\mathbf{x})$ during the optimization. The studies by Loshchilov, Schoenauer, and Sebag (2012), Marsden, Wang, Dennis Jr, and Moin (2004), Ong, Nair, Keane, and Wong (2005) and Won and Ray (2004) feature benchmark comparisons of surrogate-based optimization. Nevertheless, the model construction and updating of the surrogates also requires computational resources, as well evaluations for verification on the more expensive function layers. An advantage of surrogate-based optimization is the availability of the surrogate model, which can be utilized to gain further global insight into the problem, which is particularly valuable for black-box problems. For instance, the surrogate can be utilized to identify important decision variables or visualize the nature of the problem, i.e., the fitness landscape.

A common optimization process using surrogates is outlined by the following steps:

1. Sampling the objective function at k positions with $\mathbf{y}_i = f_1(\mathbf{x}_i)$ or $\mathbf{y}_i = f_2(\mathbf{x}_i)$, $1 \leq i \leq k$ to generate a set of observations $\mathcal{D}_t = \{(\mathbf{x}_i, \mathbf{y}_i), 1 \leq i \leq k\}$. The sampling design plan is commonly selected according to the surrogate.
2. Selecting a suitable surrogate. The selection of the correct surrogate type can be a computational demanding step in the optimization process, as often no prior information indicating the best type is available. Common types of surrogates will be presented in Section 7.

3. Constructing the surrogate $s(\mathbf{x})$ using the observations.
4. Utilizing the surrogate $s(\mathbf{x})$ to predict n new promising candidates $\{\mathbf{x}_{1:n}^*\}$, e.g., by optimization of the *infill* function with a suitable algorithm. For example, it is reasonable to use an exact algorithm, as the surrogate often provides the required global information or is very cheap to evaluate.
5. Evaluating the new candidates with the objective function $y_i^* = f_1(\mathbf{x}_i^*)$ or $y_i^* = f_2(\mathbf{x}_i^*)$, $1 \leq i \leq n$.
6. If the stopping criterium is not met:
Updating the surrogate with the new observations $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(\mathbf{x}_i^*, y_i^*), 1 \leq i \leq n\}$, and repeating the optimization cycle (4.-6.)

The accuracy of a surrogate strongly relies on the selection of the correct model type to approximate the objective function. Furthermore, it relies on the method of initialization and the connected initially available information, i.e., by sampling candidate solutions by means of appropriate designs. These initial sampling design plans have a major impact on the performance and should be carefully selected. Another important aspect is the selection of the next candidate. A number of the available surrogates provide information about the fitness distribution with mean and variance at the candidate location. This information can be used in the optimization process to apply sophisticated *infill functions* for predicting promising candidates. The most elementary *infill criterion* is the *best predicted improvement* on the surrogate. Sophisticated infill criteria include the well-known *expected improvement*, or confidence bound criteria (Mockus, 1974; Jones, Schonlau, and Welch, 1998; Schonlau, 1997). Particular expected improvement focuses on maintaining a balance between exploration and exploitation. By selecting of a certain surrogate, the user makes certain assumptions regarding the characteristics of the objective function, i.e., modality, continuity and smoothness (Forrester and Keane, 2009). Most surrogates are selected to provide continuous, low-modal and smooth landscapes, which renders the optimization process computational inexpensive and straightforward in comparison to the objective function, which often possesses unknown properties.

3 A new Intuitive Taxonomy

A taxonomy is defined as a consistent procedure or classification scheme for separating objects into classes or categories on basis of certain features. The term *taxonomy* is particular present in natural science for establishing hierarchical classifications. A taxonomy fulfills the task of distinction and ordering; it provides explanations and a greater understanding of the research area through the identification of coherences and the differences between the classes.

As a starting point, we will take a look at existing taxonomies, because many different surveys and handbooks for specialized optimization algorithms and associated techniques, such as metaheuristics (Boussaïd et al., 2013), direct search methods (Audet, 2014; Kolda, Lewis, and Torczon, 2003), or nature-inspired methods (Rozenberg, Bck, and Kok, 2011) are available. Moreover, we will investigate aspects of algorithm design and derive classification features as a basis for the different classes of algorithms. Further, we suggest common or manageable problem properties for each of the identified classes, e.g., the type of the underlying problem search space or the resource cost of function evaluations.

3.1 History of Taxonomies

This section illustrates a time line of the different CGO algorithm taxonomies found in the literature.

A first overview on global optimization was presented by Leon (1966). Leon classified algorithms into three categories: 1. *Blind search*, 2. *Local search* 3. *Non-local search*.

In this context, *blind search* refers to strategies where the candidates are selected at random in the complete search space but following a built-in sequential search strategy. During the *local search*, new candidates are selected only in the immediate neighborhood of the previous candidates. This leads to a trajectory of small steps. Finally, *non-local search* allows to escape from local optima and thus enables a global search.

Archetti and Schoen (1984) proposed a taxonomy of global optimization which for the first time includes surrogate-based approaches. Their taxonomy reads as follows:

1. Deterministic:
 - 1.1. Space covering methods
 - 1.2. Trajectory and tunneling methods
2. Probabilistic:
 - 2.1. Random sampling methods
 - 2.2. Random search methods
 - 2.3. Methods based on a stochastic model (*i.e.*, a *surrogate*) of the objective function

They used the term *deterministic* for the class of algorithms, which we classified as *exact* or *complete*, *i.e.*, those who are guaranteed to find the global optimum with a defined budget. The deterministic class encompasses space covering and more complex trajectory methods, where the candidate is following a trajectory which passes through all local minima. Additionally, following the scheme, the term *probabilistic* defines those algorithms which allow non-exact solutions, *i.e.*, heuristic or metaheuristic approaches.

Random sampling methods are by these means described as algorithms, which perform a set of local searches starting from different, uniformly sampled initial points. Random search methods iteratively alter a candidate solution utilizing random distributions. The paper stands out in establishing a taxonomy, which for the first time includes the concept to construct a surrogate model of the objective function. This is the groundbreaking idea of modern surrogate-based optimization algorithms. Surrogates are described as stochastic models based on evaluated points with successive updates in order to select proper candidates. Moreover, Gaussian field models are already mentioned as possible models for multivariate surrogates. However, no concrete methods or frameworks applying these ideas are described.

Törn and Zilinskas (1989) reviewed existing classification schemes and presented their own classifications. They stated that the most important distinction has to be made between two non-overlapping main classes, namely those methods with guaranteed accuracy and those without. The main new feature of their taxonomy is the clear separation of the heuristic methods in those with *direct* and *indirect* objective function evaluation, resulting in the following scheme of three classes:

1. Methods with guaranteed accuracy
 - 1.1. Covering methods

2. Direct Methods
 - 2.1. Random search methods
 - 2.2. Clustering methods
 - 2.3. Generalized descent methods
3. Indirect Methods:
 - 3.1. Methods approximating the level sets
 - 3.2. Methods approximating the objective function

Methods with guaranteed accuracy are similar to exact algorithms as aforementioned in Neumaier (2004). Direct methods utilize local function evaluations and encompass algorithms with pure random search, single start and multi start strategies. Furthermore, they include clustering methods, which try to identify regions of attraction of local minima to guide the search process and prevent reevaluation of known regions. Indirect methods use evaluations to build global models, similar to the taxonomy of Archetti and Schoen (1984), these methods cover the ideas of modern surrogate-based optimization. The use of Bayesian optimization (Mockus, 1974) was also discussed. Today's high availability of computational power did not exist, therefore Törn and Zilinskas (1989) concluded the following regarding Bayesian models and their applicability for (surrogate-based) optimization:

Even if it is very attractive theoretically it is too complicated for algorithmic realization. Because of the fairly cumbersome computations involving operations with the inverse of the covariance matrix and complicated auxiliary optimization problems the resort has been to use simplified models.

Moreover, Žilinskas (1992) added more insight in a specific review on statistical models for global optimization. In this work, he explicitly described common challenges in using stochastic models during in an optimization process. Special remark was given to the problems arising by using multivariate and high dimensional models, particular the need of complex computations. Thus, he suggested using simplified, reduced models.

Arora, Elwakeil, Chahande, and Hsieh (1995) presented a review of optimization algorithms with the focus on engineering applications. The paper's significant feature is the emphasis on the specific characteristics of real world problem optimization, including different constraint handling techniques. The features of different underlying problems are discussed, such as availability of function information. Moreover, they outline certain algorithm features which are needed in algorithm selection on basis of the desired optimization goal. Although the main classification is similar to the one by Archetti and Schoen (1984) and splits the algorithms into *deterministic* and *stochastic* classes, it further adds some knowledge by discussing different heuristic and metaheuristic approaches, including modern metaheuristics, such as simulated annealing (Kirkpatrick, Gelatt, Vecchi, et al., 1983), tabu search (Glover, 1989) and genetic algorithms (Eiben and Smith, 2015). The taxonomy characterized algorithms based on whether they

- are deterministic or stochastic,
- are able to solve continuous, discontinuous or combinatorial problems,
- can find all optima of a problem,
- have a local and/or global search phase,
- utilize local search, or
- need gradients.

Jones et al. (1998) were not the first ones to apply surrogate-based optimization, but significantly influenced their popularity. The *efficient global optimization* algorithm is still used as an example for a surrogate-based strategy as common in modern frameworks. It already included infill functions and a sophisticated optimization method for their optimization instead of simple multi-start or grid approaches (see Section 7.2.1). In addition, the framework distinguishes from previous work by introducing an explicit phase of model validation utilizing statistical measures.

Jones (2001) present a taxonomy of surrogate-based optimization in which he segmented the different approaches. On the surrogate model level, the distinction was drawn between

- *non-interpolating* models, which minimize the sum of squared errors for a selected functional form.
- *interpolating* models, where the function passes through all evaluated points.

Moreover, a difference between

- *Two-stage* approaches, where first a surrogate is fit to a set of preselected and evaluated candidates and then utilized for optimization and
- *One-stage* approaches, where the candidate selection is made on the basis of a hypothesis and part of the modeling process, is shown.

An extensive taxonomy focusing on exact (see Section 2.3) CGO methods was given by Neumaier (2004). The particular interesting feature of this taxonomy is that it includes a distinction between two classes:

1. *Incomplete* and *asymptotically complete* methods, which can get stuck in local optima. Furthermore, they might be able to reach the global optimum in finite time but are not able to ensure they found the global optimum.
2. *Complete* and *rigorous* methods, which reach the global minimum with certainty after a finite runtime and are able to recognize that they found it or are able to find the global minimum within given tolerances.

He also pointed out that the later class is often referred to as *deterministic*, which can be misleading as incomplete methods are also commonly deterministic.

Zlochin et al. (2004) presented a survey of model-based optimization. Although the survey focused on algorithms for solving combinatorial problems, the underlying classification is attractive as it explicitly features modern model-based optimization as a separate class. It draws a clear distinction between:

- *Instance-based* algorithms, which generate new candidates only based on the current candidate solution or candidate population and
- *Model-based* algorithms. This class encompasses algorithms which generate candidates by utilizing a parameterized probabilistic model, which is updated using evaluated points, as done in distribution-based algorithms (see Section 6.3).

The authors pointed out that *model-based* algorithms describe an adaptive stochastic mechanism which generates candidates and does not approximate the objective function. Thus, they explicitly excluded surrogate-based optimization as previously introduced in this article.

In difference to the former classifications, Boussaïd et al. (2013) focused on a concrete design aspect of optimization algorithms and divided them into two classes:

1. *Single-solution* based metaheuristics which start with a single initial solution and then make step-wise movements away from this solution, thus forming a trajectory. They also refer to this class as *trajectory* methods. According to the authors, these algorithms can be seen as *intelligent* extensions of local search algorithms.

2. *Population-based* methods, i.e., methods which use a set of concurrent candidate solutions instead a single solution.

A summary of the algorithm classifications discussed so far is given in Figure 2. For an easier comparison, we divided the algorithms roughly in three distinct classes:

1. *Exact* 2. *Heuristic and metaheuristic* 3. *Surrogate optimization*

For the sake of completeness, we added a description of our new taxonomy, which we present later in this section. The presented new taxonomy is intended to be a successor of the taxonomy presented by Törn and Zilinskas (1989), which is the last in our overview to include all classes and provide a general classification.

3.2 Motivation for a new Taxonomy

Considering the history of CGO taxonomies (Section 3.1, cf. Figure 2), we can conclude that during the last decades, several new taxonomies for these optimization algorithms were developed. However, new algorithms are proposed nearly every day and particular model-based and surrogate-based algorithms dominate the field of real-world applications and have become the state-of-the-art in modern algorithm design for this task. Existing taxonomies of CGO algorithms do not reflect this situation. Although there are surveys and books which handle the broad field of optimization and give general taxonomies, they are outdated and lack the integration of the new designs. Available up-to-date taxonomies often specialize on a subfield of algorithms and do not attempt to present a general overview, which allows to connect and compare the different optimization algorithms. This motivated the development of a new taxonomy that explicitly includes surrogate-based optimization and puts it in the larger context of CGO algorithms.

From our point of view, it is important that our new taxonomy not only divides the algorithms into classes, but also provides an understanding of the working mechanisms of each algorithm. This grants an insight for which kind of problem they are suitable. It is thus crucial, that the basic concept of the taxonomy is *comprehensible* and appears *intuitive* to a broad audience.

In his work about evolution strategies, Rechenberg (1994) illustrated a visual approach to an optimization process: a mountaineer in an alpine landscape, attempting to find and climb the largest mountain. We will further investigate the idea of optimization processes being human-like individuals trying to find their way to the most attractive location and thus define our main class names: The *Wanderer*, the *Guide*, the *Cartographer* and two additional classes, the *Hybrid* and *Exact*. This naming convention shall accomplish two goals:

1. Giving a simple idea on how the algorithms in the associated class perform their basic search.
2. Being obvious and straightforward and consequently, being simple to memorize.

To reach our first goal, we will identify the crucial elements of each optimization algorithm. These elements allow us to extract classification features, which can be used to establish a new distinct taxonomy and further accomplishes our second goal. To further support the comprehensibility of our taxonomy, we will draw an *analogy* between the algorithm class names and the behavior of a human-like individual in each of the descriptive class section. The usage of analogies to the natural world is not a new idea. Instead, it is common in the area of metaheuristics, where the behavior of animals inspires the search procedure of the algorithms: Evolutionary algorithms are based on the evolution theory (Rechenberg, 1994; Eiben and Smith, 2015);

	Exact	Heuristic and Metaheuristic	Surrogate-based Optimization
1966 Leon		Blind search, Local search, Non-local search	
1984 Archetti	Deterministic methods, Covering methods	Probabilistic methods, Random sampling, Random search methods	Stochastic model
1989 Törn and Žilinskas 1992 Žilinskas	Guaranteed accuracy, Covering methods	Direct methods, Random search methods, Clustering methods	Indirect methods, Approximating objective function
1995 Arora	Deterministic	Stochastic	
2001 Jones			Non-Interpolating, Interpolating, Two-stage, One-stage
2004 Neumaier	Complete and rigorous	Incomplete and Asymptotically Complete	
2004 Zlochin		Instance-Based, Model-Based	
2013 Boussaïd		Single-Solution, Population-Based	
2018 Stork	Exact	Wanderer, Guide	Cartographer

Figure 2: Continuous Global Optimization Taxonomy History. Information from Leon (1966), Archetti and Schoen (1984), Törn and Zilinskas (1989), Arora et al. (1995), Jones et al. (1998), Pintér (2002), Neumaier (2004), Zlochin et al. (2004) and Boussaïd et al. (2013) are illustrated and compared. Different distinctions between the large set of algorithms were drawn. The lack of a modern complete taxonomy of CGO algorithms, which comes into legacy of the taxonomy presented by Törn and Zilinskas (1989), is obvious. The diagram concludes with the taxonomy presented in Section 3.4.

particle swarm optimization (Kennedy and Eberhart, 1995; Shi and Eberhart, 1998) uses the movement of bird flocks as a role model; ant colony optimization (Dorigo, Birattari, and Stutzle, 2006) mimics, as the name suggests, the ingenious path finding and food search principles of ant populations. These two examples indicate that analogies are useful to inspire developers to create new search strategies. They are also helpful to explain the behavior of these search algorithms, which makes them valuable to be used in a comprehensible taxonomy.

3.3 The Five Elements of Algorithm Design

Any modern optimization algorithm, as defined in Section 2.2, can be reduced to the four key search strategy elements *Initialization*, *Variation*, *Evaluation* and *Selection*. All these key elements are controlled by a fifth element: the *control parameters* for the different functions and operators in each element. Algorithm 3.1 displays the key elements and the abstracted fundamental structure of optimization algorithms (Bartz-Beielstein and Zaefferer, 2017). These structure and elements can be mapped to any modern optimization algorithm. Even if the search strategy is inherently different or elements do not follow the illustrated order or appear multiple times per iteration.

Algorithm 3.1: General Optimization Algorithm

```

1 set initial control parameters
2 begin
3    $t = 0$ 
4   initialize candidate(s)
5   evaluate initial candidate(s)
6   while not termination-condition do
7      $t = t + 1$ 
8     variate solutions to get new candidate(s)
9     evaluate new candidate(s)
10    select solution(s) for next iteration
11    optional: update control parameters
12  end
13 end

```

The **initialization** of the search defines starting locations or a schema for the initial candidate solutions. Two common strategies exist:

1. If there is no a priori knowledge about the problem and its search space, the best option is to use strategically randomized starting points. Particularly interesting for surrogate-based optimization are systematic initialization schemes by methods from the field of *design of experiments*.
2. If domain knowledge or other a priori information is available, such as information from the process or data from previous optimization runs, the algorithm should be initialized utilizing this information to full extend, e.g., by using a selection of these solutions, such as these with best fitness. In surrogate-based optimization available data can be used for the initial modeling.

The initial candidates have a large impact on the balance between exploration and exploitation. Space-filling designs with large amounts of random candidates or sophisticated *design of experiments* methods will lead to a initial exploration of the search space. Starting with a single candidate will presumably lead to an exploitation of the neighborhood of the selected candidate location. Hence, algorithms using the first scheme are in general more robust, while the latter are sensitive to the selection of the starting candidate, particular in multi-modal landscapes. The robustness can be further increased by multi-start strategies, which are particular common for single-candidate algorithms, and also frequently recommended for population-based algorithms (Hansen, Auger, Ros, Finck, and Pošík, 2010a).

The **variation** during the search process defines the methods for generating new candidates, with special regard on how available or obtained information about the objective function is used. A standard approach is the variation of existing observations, as it utilizes, and to a certain extend preserves, the information of previous iterations. Even the simplest *wanderer* class algorithms (Sec 5.1), which do not require any global information or stored knowledge of former iterations, utilize the last obtained solution to generate new candidate(s). Sophisticated algorithms generate new candidates on the basis of exploited and stored global knowledge about the objective function and fitness landscape. This can be conducted explicitly by either keeping an archive of all available or selected observations, or implicitly by using distribution or data models of available observations. Another option to generate new candidates is combining information of multiple candidates by dedicated functions or operators, particular present in the guide class (Sec 6.1). The exact operators for generation and variation of candidate solutions are various and a key aspect of keeping the balance between exploration and exploitation in a search strategy.

The **evaluation** defines how the fitness of the candidates is computed and which objective function is utilized. The evaluation is the key aspect of any algorithm, as it defines the basis for any information gain and has a huge influence on the search strategy. For black-box problems, the evaluation of candidates is the only option to exploit any problem information. We differentiate between a *direct* evaluation of the objective and *indirect* evaluation by using the predicted fitness provided by a surrogate. How evaluation is performed depends mainly on the underlying problem and is largely influenced by the design of the objective function. Important aspects in real world problems are *noise*, *constraints* and *multiple objectives*.

While most computer experiments can be seen as deterministic, i.e., iterations using the same value set for the associated decision variables should deliver the same results, real-world problems are often non-deterministic. They include non-observable disturbance variables and stochastic *noise*. Typical noise handling techniques include multiple evaluations of solutions to reduce the standard deviation and special sampling techniques. The interested reader can find a survey on noise handling by Arnold and Beyer (2003).

Moreover, many real-world problems frequently include different constraints, which need to be considered during the optimization process. Constraint handling techniques can be directly part of the optimization algorithm, but most algorithms are designed to minimize the objective function and constraint handling is added on top. Thus, it is often integrated by adjusting the fitness, e.g., by penalty terms. Different techniques for constraint handling are discussed by Coello (2002) and Arnold and Hansen (2012).

The evaluation of multiple objectives can include several correlated objective functions and usually delivers a set of non-dominated solutions, a so-called Pareto-set (Naujoks, Beume, and Emmerich, 2005). In this case, a so-called decision maker is utilized to compute the fitness of a

solution and select solutions from the Pareto-set (Fonseca, Fleming, et al., 1993).

The **selection** defines the principle of choosing the solutions which will be used in the next iteration. We use the term *selection*, which has its origins in evolutionary computation. Beside the simplest strategy of choosing the solution(s) with the *best* fitness, advanced selection strategies have emerged, which are particular present in metaheuristics (Boussaïd et al., 2013). These selection strategies are particular common in algorithms with several candidates per *variation* step, thus the most sophisticated selection methods were introduced in the scope of evolutionary computation (Eiben and Smith, 2015). A common strategy is based on relative fitness comparisons, so-called *ranked* selection. Detailed examples of selection strategies are given in Section 6.2.

Control parameters determine how the search can be adapted and improved by controlling the above mentioned key elements. We distinguish between *internal* and *external* parameters: External parameters, also known as offline parameters, can be adjusted by the user and need to be set a priori to the optimization run. Typical external parameters include the number of candidates and settings influencing the above mentioned key elements. Beside common theory-based defaults (Schwefel, 1993), they are usually set by either utilizing available domain knowledge, extensive a priori benchmark experiments (Gämperle, Müller, and Koumoutsakos, 2002), or educated guessing. Sophisticated *tuning* methods were developed to exploit good parameter settings in an automated fashion. Well known examples are sequential parameter tuning (Bartz-Beielstein et al., 2005), iterated racing for automatic algorithm tuning (López-Ibáñez, Dubois-Lacoste, Cáceres, Birattari, and Stützle, 2016), *bonesa* (Smit and Eiben, 2011) or *SMAC* (Hutter, Hoos, and Leyton-Brown, 2011).

In difference, internal parameters are not meant to be changed by the user. They are either fixed to a certain value, which is usually based on physical constants or extensive testing by the authors of the algorithm, or are *self-adaptive*. Self-adaption or online control changes the parameters online during the search process on basis of the gathered knowledge or exploited problem information (Eiben et al., 1999) without user influence. Algorithms using self-adaptive schemes thus tend to gain outstanding generalization abilities and are especially interesting for black-box problems, where no information about the objective function properties is available (Hansen, Müller, and Koumoutsakos, 2003). In general, the settings of algorithm control parameters directly affect the balance between exploration and exploitation during the search and are crucial for the performance.

3.4 Features of an Intuitive Taxonomy

Taxonomies are often based on the subjective author's experience and former definitions, as well as more impartial features and similarities. Prior to establishing our new taxonomy, we defined a set five essential *classification features* (CF), which are based on former available taxonomies as well as significant features and similarities of CGO algorithms. They are intended to give a good understanding on how we separated our classes to create a distinct taxonomy, which still remains comprehensible and intuitive. Prior to each class description we will outline their features in a text box as exemplified in Figure 3. We will refer to our classification features as CF-I to CF-V:

CF I) Use of Information:

The information feature has four possible categories: The first category is *memoryless*. The term describes algorithms, which only use the available information of the prior iteration (or initialization).

The second is *explicit* memory. It defines those algorithms, which information of prior iterations in a direct fashion, e.g., by maintaining an archive of all observations.

Third is *implicit* memory; these algorithms combine information of several iterations and solutions by operators, functions or models.

Finally, algorithms which require *a priori* information about the objective function, such as the value of the optimum.

CF II) Candidate Evaluation:

The candidate evaluation feature defines if the objective function value is only *directly* calculated or also *indirectly* approximated. The approximation of the fitness during the candidate *variation* phase can greatly lower the necessary amount of objective function evaluations, but no optimization process can be reliable and successful without verification of these candidates with the objective function

CF III) Type of Candidate:

This feature refers to the number and type of candidate solutions used in the *variation* and maintained in each iteration. It has three categories:

The first is *single* and implies that the variation is based on a single candidate solution. Moreover, these algorithms maintain only a single solution for their next iteration.

The second type is *population*, where in each iteration the variation is based on several candidate solutions and moreover, several solutions are stored for the next iteration.

The most sophisticated type are *model*-based algorithms, which utilize a candidate distribution model for the variation, which is stored and adapted in each iteration. The candidate evaluation is not affected and remains *direct*.

CF IV) Region of Search:

This feature describes the effective search region of an algorithm.

Local algorithms have no operators or functions for exploration and are thus not capable of escaping the so-called *region of attraction* of an optimum.

Global algorithms have the ability to find optima in multi-modal landscapes by introducing operators or functions to balance exploration and exploitation.

CF V) Problem Properties:

This feature is a collection of objective function properties. Algorithms which are efficient in solving a problem with the given property are assigned with the related feature.

- **Domain Knowledge:** This feature describes problems with known function properties, such as a mathematical problem formulation or information about the number and objective function value of optima. This knowledge can be exploited and used for a efficient or exact search process.
- **Unimodal:** This term describes objective functions with a single optimum in a linear or convex search space.

- **Multimodal:** Problems are called multimodal when they have several local and/or global optima.
- **Black-Box:** Problems are called black-box if they do not provide any domain knowledge and all information needs to be gathered by objective function evaluations. Many real-world problems can be associated with this characteristic.
- **Discontinuities:** Functions who have jumps in their objective function values are discontinuous and not differentiable.
- **Noisy:** Noisy objective function do not return deterministic function values. Multiple evaluations of the same candidate on a noisy function can lead to different results.
- **Multi-Objective:** Multi-Objective problems have several (competing) objectives.
- **Expensive:** Expensive problems have a high cost for each function evaluation in terms of either physical resources or computation time.

Classification Features:

- I Use of Information: *a priori, memoryless, explicit/implicit memory*
- II Candidate Evaluation: *direct, indirect*
- III Type of Candidate: *single, population, model*
- IV Search Space: *local, global*
- V Problem Properties: *unimodal, multimodal, black-box, ...*

Figure 3: Fundamental features for algorithm classification

For an overview of all classes and the connected classification features *CF I-IV*, we utilize a decision tree, illustrated in Figure 4. In this figure, we utilized the features I to IV as classification nodes which conclude in our new taxonomy with the introduced classes *exact, wanderer, guide* and *cartographer*. Moreover, the associated *problem properties* are displayed. The figure is intended to provide a fast overview on how the new taxonomy works. New algorithms could be easily included by assigning the above listed features and then using the illustrated decision tree. We know that not all algorithms will fit into our taxonomy, as they have property combinations, which are not displayed in our scheme. They may also belong to the *hybrid* class, which represent combinations of methods from the displayed classes. The hybrid class is not shown in the figure, as it has no distinct properties.

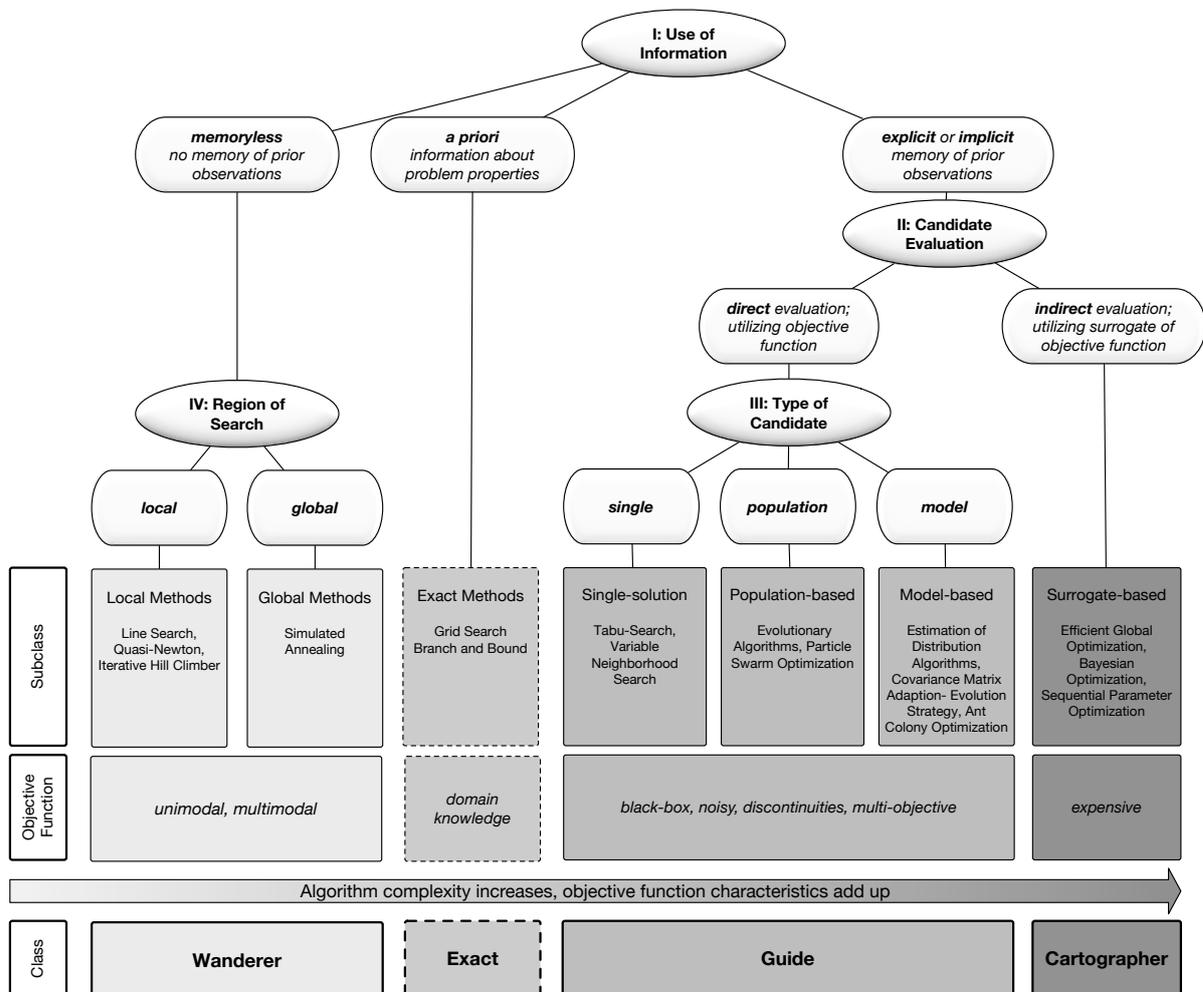


Figure 4: Algorithmic classification including the classes exact, wanderer, guide, and cartographer. The defined classification features I to IV are used as nodes of a decision tree with the subclasses and related example algorithms as the final leaves. Below the subclass the corresponding main classes and the associated objective function characteristics are illustrated. The objective function characteristics are displayed bottom-up, i.e., the displayed characteristics are adding up and getting more sophisticated from left to right.

4 The Exact Class

Exact Class:

- I Use of Information: *a priori*
- II Candidate Evaluation: *direct*
- III Type of Candidate: *single*
- IV Search Space: *global*
- V Problem Properties: *unimodal, multimodal, domain knowledge*

The general description of exact algorithms was presented in Section 2.3. For an efficient search, exact algorithms need a priori information about the objective function. Therefore, they are only suitable for a limited class of problems where this information is available. Moreover, the application of exact algorithms is always a tradeoff decision between computation time and precision.

For example: Given a nondeterministic polynomial (NP)-hard problem, usually no exact algorithm exists that is able to find the best solution in polynomial time. The traveling salesman problem is a common combinatorial NP-hard problem, where the goal is to find the minimal length tour through a fixed number of cities. An exact algorithm could solve this problem by calculating every possible tour and selecting the best. Despite providing the best solutions, this strategy would use a lot of computation time. Particular if the number of cities gets large, this problem cannot be solved exact in reasonable time.

It should be noted that even if exact algorithms are not efficient for NP-hard and black-box problems, they are commonly part of surrogate-based optimization frameworks, because the models provide the required information for an exact search or are computationally very cheap to evaluate (see Section 2.4 and Figure 1). For example, *efficient global optimization* (Section 7.2.1) and *bayesian optimization* (Section 7.2.2) can use exact algorithms.

Two common methods from the family of exact methods are *grid search* and *branch-and-bound*. Grid search combines a multi-start local optimization with an increasingly subtler sampling grid of starting points. Branch-and-bound optimization is conducted by splitting the original problem recursively into subproblems with the goal of excluding or solving them, until it is guaranteed that no subproblem can lead to a better solution. It is known as branch-and-bound (Lawler and Wood, 1966) as lower bounds on the objective function are computed. More complex exact search methods combine branching and local optimization, Lipschitzian optimization, convexity or interval analysis (Neumaier, 2004; Floudas, 2013; Hansen and Walster, 2003; Horst and Tuy, 2013).

Example 4.1 (Dividing Rectangles). An example for an exact algorithm is dividing rectangles (DIRECT), initially proposed by Jones, Perttunen, and Stuckman (1993) as a modification of Lipschitzian optimization. While assuming that the objective function is Lipschitz continuous, the algorithm does not need a specification of the Lipschitz constant, as it is sampled during the optimization run. The algorithm uses hypercubes to divide the search space. The center c_i of each hypercube is sampled and given a fitness, based on the objective function value and the size of the associated hypercube. Based on this fitness, the hypercube which is most likely to inherit the optimum is selected and further divided into smaller hypercubes. Then their fitness is sampled and the process repeats, until a stopping criterion is met or the algorithm has converged.

5 The Wanderer Class

The wanderer class encompasses algorithms with non-complex search strategies which generate and maintain a single candidate per iteration. New candidates are generated in the vicinity of the current solution by a stochastic process which is independent of previous search steps. The consecutive candidates describe a trajectory in the search space that forms in the ideal case a direct line to the optimum. In their search, they only use the local objective function information about the prior solution. Thus, these algorithms do not use global information about the problem in the *variation* or *selection* steps.

Analogy 1 (The Wanderer).

The intuitive description of a wanderer is a single individual who *wanders* through the landscape to find the most attractive place in a given area. During its search it only utilizes local information about its current position to find the best direction. If the goal of this individual is to find the highest mountain, it will likely follow the ascending way, because it directly satisfies the current objective. It does not memorize gathered information, so that there is a chance that it will circle a position, revisit a place or gets completely lost.

We separate between the *local* and *global* subclass of the wanderer: While the local wanderer keeps a greedy *selection*, the global wanderer also allows the acceptance of non-improving candidates.

5.1 Local Wanderer

Local Wanderer:

- I Use of Information: *memoryless*
- II Candidate Evaluation: *direct*
- III Type of Candidate: *single*
- IV Search Space: *local*
- V Problem Properties: *unimodal*

The local wanderer subclass consists of basic local optimization algorithms, which include classical gradient-based algorithms as well as deterministic or stochastic hill-climbing algorithms. These algorithms are designed for fast convergence to a local optimum situated in a region of attraction $\mathcal{A} \subseteq \mathcal{S}$ and have no explicit strategy for exploration.

Gradient-based methods, such as quasi-Newton Methods (Shanno, 1970), directly compute or approximate the gradients of the objective function to find the steepest direction of the optimum. Direct-Search methods perform an iterative and gradient-free search by using a minimal amount of information about the objective function. Overviews of direct search methods were presented by Lewis, Torczon, and Trosset (2000) and Kolda et al. (2003). Moreover, the (1+1) evolution strategy with a basic *selection* operator (further explained in Section 6.2) can be associated with this class.

Example 5.1 (Iterated Stochastic Hill-Climber). The *iterated stochastic hill-climber* (Michalewicz and Fogel, 2013) is a typical example of the local wanderer subclass. It has an elementary algorithm design, where the search *variation* is stochastic and the *selection* typically greedy. In each iteration of the algorithm a new candidate \mathbf{x}_t is created by sampling from a probability distribution \mathcal{D} around the prior solution \mathbf{x}_{t-1} . The variance of the distribution, which is often

uniform or Gaussian, defines the so-called *step size* of the variation and is the most important control parameter of this algorithm. The greedy *selection* works as follows: If the new candidate \mathbf{x}_t has a better fitness value than the prior solution \mathbf{x}_{t-1} , it is accepted as new solution and the iteration is repeated. If the fitness is not improved, the prior solution is kept for the next iteration. This greedy selection scheme is based on a comparative, ranking-based selection of the candidates with no influence of the absolute difference in the objective function values. This implies invariance to linear transformations and scaling of the objective function. As the basic scheme includes no operator to escape from a local optimum or exploration, this algorithm will most likely converge to a local optimum. The termination condition is usually set to a number of non-improving iterations, after which is assumed that an optimum has been found. The generalized design of the iterated hill-climber is outlined in Algorithm 5.1.

Algorithm 5.1: Iterated Hill-Climber

```

1 begin
2    $t = 0$ 
3   initialize  $\mathbf{x}_t$  with (random) candidate  $\mathbf{x} \in \mathcal{S}$ 
4   evaluate  $\mathbf{x}_t$ 
5   while not termination-condition do
6      $t = t + 1$ 
7     sample new candidate  $\mathbf{x}_t$  from probability distribution  $\mathcal{D}$  around the current
       solution  $\mathbf{x}_{t-1}$ 
8     evaluate candidate  $\mathbf{x}_t$ 
9     if new solution improves the fitness  $\mathbf{y}_t < \mathbf{y}_{t-1}$  then
10      | accept new solution
11      | else
12      | keep old solution  $\mathbf{x}_t = \mathbf{x}_{t-1}$ 
13      | adjust variation step size/variance of probability distribution (optional)
14    end
15 end

```

□

As the name implies, the local wanderer is in the first place suitable for unimodal functions or to exploit local optima. It can be applied for global optimization to multimodal landscapes if an adequate multi-start strategy is used. These multi-start strategies typically demand a high number of function evaluations and are only reasonable to be used for problems with relatively cheap objective functions, such as in surrogate-based optimization. Theoretically, a local wanderer could establish a global search by sampling the candidates from a distribution with a dispersion over the complete search space, which is equivalent to a very large step size for the *variation*. This strategy would be inefficient as it leads to random sampling of candidates without considering any local information and direction of search. Thus, it is common to limit the dispersion of the probability distribution to the vicinity of the current solution, to form a small neighborhood which is very small compared to the complete search space. This leads to the outlined *hill-climbing* search strategy which performs a trajectory of small, fitness-improving steps. The maximal step size is consequently an important control parameter, which is often designed as being adaptive in many versions of the algorithm (cf. mutation operator in Section

6.2). In general, local wanderers are often part of sophisticated algorithms as fast converging local search strategy.

5.2 Global Wanderer

Global Wanderer:

- I Use of Information: *memoryless*
- II Candidate Evaluation: *direct*
- III Type of Candidate: *single*
- IV Search Space: *global*
- V Problem Properties: *multimodal*

The global wanderer subclass encompasses algorithms which implement operators to balance exploration and exploitation in the *selection* process of candidates. They differ from local wanderers by their explorative search strategies, which further enable global optimization. Similar to the local version, the global wanderer utilizes a stochastic *variation* of candidates in the neighborhood of the current solution without considering stored or modeled global information. Exploration is achieved by introducing operators or functions which allow to expand the search space and escape the *region of attraction* of a local optimum. The well-known *simulated annealing* (SANN) will be used for exemplifying this approach.

Example 5.2 (Simulated Annealing). Kirkpatrick et al. (1983) introduced SANN as a search procedure for global combinatorial optimization. It is known to be a significant contribution to the field of metaheuristic search algorithms. The basic search procedure of every SANN method is inspired by annealing in metallurgy, where a material is first heated, then cooled. Following this analogy, the most important element to control the search is an adaptive parameter: the *temperature* (T). Analogous to the thermodynamic energy in a heated material, it increases the possible movement of the candidates.

The continuous version (Goffe, Ferrier, and Rogers, 1994; Siarry, Berthiau, Durdin, and Haussy, 1997; Van Groenigen and Stein, 1998) of the SANN algorithm basically extends the iterated stochastic hill climber. It includes a new element to allow global search, the so-called acceptance function $P(\mathbf{x}_{t-1}, \mathbf{x}_t, T)$. The acceptance function is used during the *selection* and determines the probability of accepting an inferior candidate as solution by utilizing T as a parameter. This dynamic *selection* allows to escape local optima steps by accepting movement in the opposite direction of improvement, which is the fundamental difference to a hill climber. A common example for an acceptance function is the so-called *Metropolis* function:

$$P(\mathbf{x}_{t-1}, \mathbf{x}_t, T) = \min \left\{ 1, \exp \left\{ -\frac{f(\mathbf{x}_t) - f(\mathbf{x}_{t-1})}{T} \right\} \right\} \quad (1)$$

The Metropolis function always accepts fitness improving steps towards the minimum ($f(\mathbf{x}_t) - f(\mathbf{x}_{t-1}) \leq 0$) and moreover, has a probability to accept ascending ($f(\mathbf{x}_t) - f(\mathbf{x}_{t-1}) > 0$) steps based on T. Higher T values thus increase the probability to accept an inferior candidate. In contrast to rank-based selection schemes, this function is scale-based and utilizes absolute differences in the fitness, which renders it sensitive to linear transformations of the objective function, e.g., a multiplication by a scalar.

At the end of each iteration, a so-called *cooling* operator C adapts T. This operator can be used to balance exploration and exploitation (Henderson, Jacobson, and Johnson, 2003). A common

approach is to start with a high temperature and steadily reduce T according to the number of iterations or to utilize an exponential decay of T . This steady reduction of T leads to a phase of global exploring in the early iterations, while with decreasing T the probability of accepting inferior candidates is reduced. With approaching a T value of zero, the behavior of the algorithm becomes more and more similar to this of an iterative hill-climber. A general SANN version is displayed in Algorithm 5.2. Modern SANN implementations integrate *self-adaptive* cooling-schemes which use alternating phases of cooling and reheating (Locatelli, 2002). These allow a more sophisticated control of exploration and exploitation. The basic version of the algorithm does not use any archive, so that evaluated solutions (including exploited optima) are not stored. As the algorithm is able to escape the region of attraction of a local optima by accepting inferior solutions, the algorithm can convergence to a local optima which has an inferior fitness than the best found solution during the search. Thus, later versions commonly store at least the best found solution.

Algorithm 5.2: Simulated Annealing

```

1 begin
2    $t = 0$ 
3   initialize  $T_{t=1} > 0$ ,  $\mathbf{x}_t$  with (random) candidate  $\mathbf{x} \in \mathcal{S}$ 
4   evaluate  $x_t$ 
5   while not termination-condition do
6      $t = t + 1$ 
7     sample new candidate  $\mathbf{x}_t$  from probability distribution  $\mathcal{D}$ 
8     evaluate candidate  $\mathbf{x}_t$ 
9     if new solution satisfies acceptance function  $P(\mathbf{x}_{t-1}, \mathbf{x}_t, T_t)$  then
10    | accept new solution
11    | else
12    |   keep old solution  $\mathbf{x}_t = \mathbf{x}_{t-1}$ 
13    | vary  $T_t$  using cooling scheme  $C$  to get  $T_{t+1}$ 
14  end
15 end

```

□

Global wanderers are suitable for searches in unimodal and multimodal problems, particular if multi-start strategies are used. As they do not rely on stored information of former iterations during their search, they are also a good choice to handle dynamic objective functions (Carson and Maria, 1997; Corana, Marchesi, Martini, and Ridella, 1987; Faber, Jockenhövel, and Tsatsaronis, 2005). However, the rather simplistic utilization of exploited global information renders them not efficient for challenging and expensive optimization problems. Moreover, the control parameters have a significant effect on the performance of these algorithms and should be tuned in an offline or online fashion.

6 The Guide Class

The guide class encompasses algorithms which utilize information from consecutive function evaluations. *Guiding* is here the process of finding a direction of improvement. In comparison to the *wanderer* the guide stores information and utilizes it during the search. Moreover, it can share information and group with other guides to organize a structured crowd search.

Analogy 2 (The Guide).

The intuitive idea of this class is a single or group of traveller(s) looking for an interesting place. They try to memorize their own route, follow travel signs about interesting or forbidden paths and ask their fellow travelers to share their knowledge and give directions. Furthermore, they can be able to consolidate all gathered information and utilize it to their benefit.

We divide this class in three sub-classes, which follow different schemes to process information and use it during the search:

- (i) The *single-solution guide* (Section 6.1) utilizes an intelligent landscape structuring and partitioning process using the information of prior search iterations.
- (ii) The *population-based guide* (Section 6.2) utilizes multiple candidates in each iteration and has special operators to combine information and to select the most successful candidates.
- (iii) The *model-based guide* (Section 6.3) utilizes mathematical or statistical models to store and process the exploited information of several iterations and/or candidate evaluations.

6.1 Single-Solution Guide

Single-solution Guide:

- I Use of Information: *implicit memory*
- II Candidate Evaluation: *direct*
- III Type of Candidate: *single*
- IV Search Space: *global*
- V Problem Properties: domain knowledge, multimodal

Single-solution guides are the connecting link between the wanderer class and population-based guides. While these algorithms also allow the sampling of several solutions in one iteration, they are still based on the principle of maintaining a single solution, which perform a trajectory in the search space over consecutive iterations. These methods are also known as *trajectory methods* (Boussaïd et al., 2013) (cf. wanderer class 5.1). In difference to the wanderer class, they explicitly use operators to store und utilize information of former iterations and to guide the search process in certain directions. They are a step towards population-based algorithms, but miss the idea of explicitly using population-based operators in the *variation* and *selection* steps.

This class encompasses search space partitioning algorithms which use the exposed knowledge of former iterations. They create sub-spaces which are *forbidden* and not considered in the current search iteration, or *attractive* sub-spaces, where the search is focused on. This search space partitions are used to find a balance between exploration and exploitation, as new candidates

are placed in promising or previous unexplored parts of the search space. For example, *variable neighborhood search* (VNS) (Hansen and Mladenovic, 2003; Hansen, Mladenović, and Pérez, 2010b; Mladenović, Dražić, Kovačević-Vujčić, and Čangalović, 2008) utilizes so-called neighborhood structures, which are pre-defined local sub-spaces. The search strategy of VNS is to perform sequential local searches in these sub-spaces to exploit their local optima. The idea behind this search strategy is that by using an adequate set of sub-spaces, the chance of exploiting a local optimum which is near the global optimum increases.

Another, well-known algorithm and outstanding paradigm for this class is *Tabu Search*:

Example 6.1 (Tabu Search). Tabu search (TS) (Glover, 1989) was introduced as an optimization method for combinatorial search spaces, where the number of possible solutions is limited. The last successful candidates are put to a so-called tabu list \mathcal{T} , which defines a sub-space of all evaluated solutions. All members of the tabu list are forbidden as candidates in the current search. This process shall ensure to move away from available solutions and prevents cycling identical candidates, including local optima. The size of the tabu list is limited by a control parameter. If the limitation is reached, the oldest solution on the list is deleted and again accepted as potential candidate.

A continuous version of the TS algorithm was presented by Siarry and Berthiau (1997). In continuous search spaces, the amount of candidates in the neighborhood of a solution is only limited to the underlying numerical accuracy, which leads to a nearly infinite number of candidates. The continuous TS algorithm thus needs to create limited sub-spaces. One approach for implementing these sub-spaces are *ball-shaped* regions, first presented by Hu (1992). A ball-shaped region (or hypersphere) $B(\mathbf{x}, h)$, is defined in the neighborhood of a solution \mathbf{x}_t . It is limited by a radius h , resulting in a new sub-space $X_B = \{(\mathbf{x}_t, \mathbf{x}) \in \mathcal{S} \mid \|\mathbf{x}_t - \mathbf{x}\| \leq h\}$. If a set of $k + 1$ of these concentric regions with different radii $\{h_{0:k}\}$ is generated, they form so-called *crowns* $C_{i:k}$, defined by:

$$C_i(\mathbf{x}_t, h_i, h_{i+1}) = \{X_C \mid h_i \leq \|\mathbf{x}_t - \mathbf{x}\| < h_{i+1}\} \text{ with } i = 0 : k \quad (2)$$

Similar to the last solution in the combinatorial TS, the region C_0 around the prior solution \mathbf{x}_{t-1} is added to the tabu list \mathcal{T} . New candidates are sampled uniformly random in each of the remaining crowns $C_{1:k}$. Siarry and Berthiau (1997) also suggest a more directed search variant by selecting candidates in direction of the gradient. From the set of new candidates, the one with the best fitness is selected for the next iteration based on a *rank-based* comparison. Exploration and exploitation can be controlled by the definition of the different radii and the corresponding partitions. An example *Linear partitioning* gives the outer crowns a larger volume and can shift the focus to a more localized search, as the density of sampled candidates will decrease. The continuous TS is displayed in algorithm 6.1.

An enhanced version of the continuous TS was presented by Chelouah and Siarry (2000). Instead of hyperspheres, they use easier to handle hyperrectangles for the partition of the search space. Moreover, they define distinct alternating search phases of exploration and exploitation. \square

The main concept of the single-solution guide is to use the information of evaluated solutions and to direct the search to former unknown regions to avoid early convergence to a non-global optimum. The strategic use of sub-spaces allows a systematic control of exploration and exploitation and particularly ensures a high level of exploration. They include a large number of parameters, such as the number or size of sub-spaces, which makes them very vulnerable

Algorithm 6.1: Continuous Tabu Search

```
1 begin
2    $t = 0$ 
3   initialize  $length(\mathcal{T})$ , number of regions  $k$ ,  $x_{t,0}$  with (random) solution  $\mathbf{x} \in \mathcal{S}$ 
4   evaluate  $x_{t,0}$ 
5   while not termination-condition do
6      $t = t + 1$ 
7     create inner tabu region  $C_0$  around  $\mathbf{x}_{t-1,0}$  and add it to tabu list  $\mathcal{T}$ 
8     create outer search regions  $C_i$  around  $\mathbf{x}_{t-1,0}$  (Eq. 2)
9     sample  $k$  new candidates  $x_{t,i:k}$  uniformly in each region  $C_{1:k}$ 
10    evaluate candidates  $x_{t,1:k}$ 
11    select solution with best fitness as current solution  $x_{t,0}$ 
12    if list  $\mathcal{T}$  has reached its maximum length then
13      delete oldest tabu region from  $\mathcal{T}$ 
14    end
15 end
```

to false setups and less good universal solvers. If correctly tuned, this class algorithms are suitable and efficient for multimodal problems. Algorithms using a pre-defined separation of the search space, such as VNS, can utilize domain knowledge for the initial definition of the sub-spaces. This renders them useful for problems where the region of the global optimum is roughly known, but not particular suitable for black-box problems.

6.2 Population-based Guide

Population-based Guide:

- I Use of Information: *implicit/explicit memory* by candidate population / solution archive
- II Candidate Evaluation: *direct*
- III Type of Candidate: *population*
- IV Search Space: *global*
- V Problem Properties: multimodal, black-box, noisy, discontinuities, multi-objective

The class of population-based guides encompasses algorithms focusing on utilizing multiple candidates during the search. They have specific, population-based operators to combine gathered information of available solutions and to compute a search direction. This class comprehends population- (or swarm-) based approaches which are commonly based on processes found in the nature. Well-known examples are *particle swarm optimization (PSO)* (Kennedy and Eberhart, 1995; Shi and Eberhart, 1998) and *evolutionary algorithms (EA)*.

PSO is a metaheuristic based on swarming theories from bird flocking behavior. Birds in a swarm adapt their movement by moving in the same direction as near flying birds. This leads to a similar movement of the whole swarm. The coordinated movement helps to avoid collisions and to quickly escape from predators. Further, birds communicate information about interesting locations such as food sources or places to sit and hide. These strategies are adapted by the PSO algorithm, where the *particles* or candidate's movement is adjusted by the global

search direction of all candidates while keeping up an individual search for an optimum. This strategy results in a balanced search with explorative elements and local convergence.

We regard EAs as the state of the art in population-based optimization, as their search concepts are dominating for this field. Nearly all other population-based algorithms use similar concepts and are frequently associated with EAs. Fleming and Purshouse (2002) go as far to state:

In general, any iterative, population-based approach that uses selection and random variation to generate new solutions can be regarded as an EA.

We will thus focus on outlining the general concepts of EAs with specific instances of the utilized search strategies, methods and operators.

Example 6.2 (Evolutionary Algorithms). Evolutionary algorithms are based on the idea of evolution, reproduction and the natural selection concept of *survival of the fittest*. In general, the field of EAs goes back to four distinct developments, *evolution strategies* (ES) (Rechenberg, 1973; Schwefel, 1977), *evolutionary programming* (Fogel, Owens, and Walsh, 1966), *genetic algorithms* (Holland, 1992), and *genetic programming* (Koza, 1992). The naming of the methods and operators matches with their counterparts from biology: candidates are *individuals* which can be *selected* to take the role of *parents*, *mate* and *recombine* to give birth to *offspring*. The *genetic* information of the parents is inherited to the offspring and can also be further changed by random *mutation*. The population of individuals is *evolved* (varied, evaluated and selected) over several iterations, so-called generations, to improve the solutions.

In general, a large number of different implementations and configurations of the mentioned operators to manipulate or alter the genetic information exist (Eiben and Smith, 2015). A basic EA is outlined in Algorithm 6.2, illustrating the four common base operators of an EA. Several examples for EAs can be found in the literature, such as the well-known $\mu + \lambda$ *evolution strategies* (Schwefel, 1993), which we will later use to exemplify some *selection* operators.

Algorithm 6.2: Evolutionary Algorithm

```

1 begin
2    $t = 1$ 
3   initialize  $\mu$  random population  $\mathcal{P}_t = \{\mathbf{x}_{t,1:\mu}\} \subseteq \mathcal{S}$ 
4   evaluate  $\mu$  random population  $\mathcal{P}_t$ 
5   while not termination-condition do
6     select parents from population  $\mathcal{P}_t^* \subseteq \mathcal{P}_t$ 
7     recombine selected parents  $\mathcal{P}_t^*$  to create  $\lambda$  offspring candidates
8      $\mathcal{O}_t = \{\mathbf{x}_{t,i}, \mu < i \leq \mu + \lambda\}$ 
9     mutate offspring
10    evaluate offspring
11    select survivors from  $\mathcal{P}_t \cup \mathcal{O}_t$  for next generation  $\mathcal{P}_{t+1}$ 
12     $t = t + 1$ 
13  end

```

Parent selection is the first of two *selection* operators in an EA. It is the process of choosing solutions from the population with size $\mu \in \mathbb{N}$, which will be utilized in the *recombination* and *mutation* step to create new candidates, the so-called offspring population with size $\lambda \in \mathbb{N}$.

The most common and simplest strategy is linear *rank-based* selection: Based on fitness comparisons, only the highest ranked candidates are selected to reproduce, leading to a high *selection pressure*. The *selection pressure* defines how likely an inferior candidate can be successful in the selection. The usage of probabilistic selection strategies which have a more dynamic chance of selection with less selection pressure is more common.

For instance, in *roulette* wheel selection, the chance of being selected is proportional to the ranking while all chances sum up to one. Each parent is then chosen by a spin of the roulette wheel, where the highest ranked individual has the highest chance of being selected.

tournament selection draws randomly small subsets of the population, commonly two or up to $\mu - 1$, for a number of tournaments. Within these tournaments, the best solution(s) are selected on basis of direct fitness comparisons.

Another, rather unusual, scheme is *fitness proportional selection*: The probability to select a solution depends on the objective function value compared to the mean objective value of the population. The disadvantage of this schema is the usage of absolute fitness, which can result in a high selection pressure and the best solution dominating the population, ultimately leading to premature convergence. This is also true for other *scale-based* selection strategies.

A complete random selection can also be used, e.g., by *uniform selection*, where each solution has the same chance of being selected, regardless of their fitness or objective function value.

Recombination is an outstanding concept of EAs and an *variation* operator. New candidates are generated by combining the information of two (or more) evaluated solutions, based on the idea of the selected, commonly fittest, parents inheriting their children their genetic information. A typical recombination operator is the *crossover*, where parts of each parent are swapped, i.e., the decision variable settings, also known as *discrete recombination*. For example, in a one-point crossover of between two discrete, string-based parents $aa|aaa$ and $bb|bbb$, two children $aa|bbb$ and $bb|aaa$ would be created, where “|” marks the crossover point. The crossover point is typically chosen at random. For real-valued variables in CGO, a basic recombination is the *arithmetic* or *intermediate recombination*, where the values are not swapped, but a value between the selected decision variables of the parents \mathbf{x}_a and \mathbf{x}_b is chosen by $\mathbf{x}_{offspring} = \alpha \times \mathbf{x}_a + (1 - \alpha) \times \mathbf{x}_b$. The α value in range $[0, 1]$ is either chosen at random, or a constant control parameter, the *crossover strength*. By $\alpha = 0.5$, *uniform arithmetic recombination* is applied. This basic recombination therefore creates candidates which are situated between both parents.

Mutation is applied to λ offspring candidates as another *variation* operator to add new genetic information and allow a higher diversification of the offspring population. Mutation is the process of changing a single candidate, commonly performed by adding a random valued vector, sampled from a parametrized uniform or Gaussian distribution. Hereby, the state of the art is to use an adaptive *mutation strength* (variation step size). It is then often set on basis of previous successful steps, for example as defined in the famous *1/5 rule* (Rechenberg, 1973). Moreover, strategies exist, which implement the mutation strength as part of the genetic information of an individual, which also undergoes variation and selection and is thus *self-adaptive*. Different concepts for self-adaptive mutation strategies can be found in the publication by Back (1996).

Survivor selection is the last step of each iteration and selects the candidates for the next iteration. While the survivor selection can generally be performed by the same selection operators as the parent selection, some special concepts were introduced and are commonly used. Two well-known concepts arise from the Evolution Strategies (Schwefel, 1993), known as $(\mu + \lambda)$ and (μ, λ) selection. The $(\mu + \lambda)$ -strategy selects survivors from the merged set of μ parents

and λ offspring candidates accordingly to their ranked fitness, whereby the top μ solutions are kept for the next generation. In the (μ, λ) -strategy all parents μ are discarded and the survivors are selected only from the offspring λ according to their ranked fitness. This requires to create $\lambda \geq \mu$ candidates in the recombination step to prevent extinction of the population. Further, in *age*-based selection strategies, each solution only survives a defined number of iterations before it gets discarded. This is often combined with fitness-based selection, as in the (μ, λ) -strategy, where each solution survives only one iteration. Aging adds a handicap to *old* dominating solutions which survived several generations. This depicts an explorative strategy to escape local optima and increase the diversity of the population. \square

EAs are very flexible in their implementation and adaptable by tuning. They are robust and suitable to solve a large class of problems, including multimodal, multi-objective, dynamic and black-box problems, even with noise or discontinuities in the fitness landscape (Jin and Branke, 2005; Marler and Arora, 2004). Further, they have successfully been applied to a large number of different industrial problems (Fleming and Purshouse, 2002) but typically require a relatively large number of function evaluations to converge. This makes them not the first choice for expensive problems where the number of evaluations is strongly limited.

The flexibility and robustness of EAs is caused by different mechanisms and strategies for controlling the balance between exploration and exploitation. A good overview is presented in the survey by Črepinšek, Liu, and Mernik (2013). The detailed survey classifies the different available evolutionary approaches and presents an intensive discussion which mechanisms influence exploration and exploitation. Theoretical aspects of evolutionary operators are discussed by Beyer (2013). The performance of EAs is influenced by *parameter tuning and control*, e.g., the setting of population size, mutation strength and selection probability. An extensive overview of the different on- and offline tuning approaches for parameter control in EAs was published by Eiben et al. (1999).

Further common strategies of controlling exploration and exploitation and multimodal optimization are so-called niching strategies, which utilize sub-populations to maintain the diversity of the population, investigate several regions of the search space in parallel or conduct defined tasks of exploring and exploiting (Shir and Bäck, 2005; Filipiak and Lipinski, 2014).

6.3 Model-based Guide

Model-based Guide

- I Use of Information: *implicit memory* by the model/distribution
- II Candidate Evaluation: *direct*
- III Type of Candidate: *model/distribution*
- IV Search Space: *global*
- V Problem Properties: multimodal, black-box, noisy, discontinuities, multi-objective

The model-based guide class encompasses algorithms, which explicitly utilize mathematical or statistical models. The distinction to surrogate-based optimization is based on the assumption, that these models do not depict direct approximations of the underlying objective function and do not target to model the complete fitness landscape. Instead, they are specialized local models which are used in the *variation*. They are not explicitly searched for optima, although often the mean of the distribution is utilized as predicted optimal solution.

A common class of model-based algorithms are *estimation of distribution algorithms* (EDA) (Larrañaga and Lozano, 2001). They generally belong to the large field of EAs (Section 6.2). The main difference to EAs is that the *variation* operators, such as recombination or mutation, are not directly applied to the candidates, but to distribution models. The distribution models are build using information of prior evaluated populations. Different distribution models can be utilized, e.g., Bayesian networks or multivariate Gaussian distributions. Further common examples of model-based algorithms are the *covariance matrix adaption - evolution strategy* (CMA-ES) (Hansen et al., 2003) and *ant colony optimization for continuous domains* $ACO_{\mathbb{R}}$ (Dorigo et al., 2006; Socha and Dorigo, 2008). $ACO_{\mathbb{R}}$ is based on the social behavior of ant colonies and their communication via pheromones. Ants search for food around their nest in a random manner and leave trails of pheromones on their way to mark and enable other ants to follow these trails. The idea of these pheromone trails are transferred to use special distribution models and *variation* operators. As the general idea of model-based algorithms is similar, we will outline their mechanisms on the example of a generalized EDA and give further insight on the particularities of $ACO_{\mathbb{R}}$ and the CMA-ES.

Example 6.3 (Estimation of Distribution Algorithms). The general idea behind the distribution model is that it is beneficial to learn the structural information of the underlying population (Larrañaga and Lozano, 2001; Hauschild and Pelikan, 2011). The structural information allows to acquire quality knowledge about the dependences among the variables. Moreover, this information is used to generate new candidates and thus, to guide the search for the optimum. A general EDA is outlined in Algorithm 6.3.

Algorithm 6.3: Estimation of Distribution Algorithm

```

1 begin
2    $t = 1$ 
3   initialize  $\mu$  random population  $\mathcal{P}_t = \{\mathbf{x}_{t,1:\mu}\} \subseteq \mathcal{S}$ 
4   evaluate  $\mu$  random candidates  $\mathcal{P}$ 
5   while not termination-condition do
6     select parents from population  $\mathcal{P}_t^* \subseteq \mathcal{P}_t$ 
7     estimate distribution parameters of parents  $\mathcal{P}_t^*$ , such as mean, variance and
      covariance
8     mutate or adapt parameters of distribution  $G$ 
9     sample distribution function  $G$  to get  $\lambda$  offspring  $\mathcal{O} = \{x_{t,1:\lambda}^*\}$ 
10    evaluate offspring
11    select survivors from  $\mathcal{P}_t \cup \mathcal{O}_t$  for next generation  $\mathcal{P}_{t+1}$ 
12     $t = t + 1$ 
13  end
14 end

```

The algorithm starts with population of μ solutions, either initialized at random or based on exposed problem knowledge. Similar to an EA, parents are selected, typically based on ranking selection. It is common to use the fittest solutions as parents (cf. selection in EAs, Sec 6.2). Instead of a direct recombination or mutation, the parents are used to construct a distribution model. For example, typical parameters of a multivariate Gaussian distribution, such as mean, variance and covariance of the selected parent population are computed. In $ACO_{\mathbb{R}}$ Gaussian

kernel functions for each dimension of the decision variable space are constructed. The distribution parameters are then target of the *variation* operators. This *variation* is the crucial aspect of the search strategy, similar to the recombination and mutation in the basic EA. For example, the CMA-ES adapts the parameters during each iteration following the history of prior successful iterations, the so-called *evolution paths*. These evolution paths are basically exponentially smoothed sums for each distribution parameter over the consecutive prior iterative steps (cf. adaptive mutation strength in Section 6.2). They thus utilize the information of several successful search steps, which is intended to quickly approach an optimum.

In the next step, a number of λ offspring candidates are generated by sampling the adapted distribution. Afterwards the candidates are evaluated and survivors for the next generation are selected, again based on survivor selection operators, as explained in (Section 6.2). The CMA-ES uses only *rank-based* selection schemes, which makes it insensitive to scaling of the objective function. \square

In general, model-based guides try to combine the benefits of statistical models and their capability of storing and processing information with population-based search operators. They are high level metaheuristics and advanced EAs which intended to be flexible, robust and applicable to a large class of problems, particular those with unknown function properties. This makes them very successful in popular black-box benchmarks (Hansen et al., 2010a). For example, the design of the CMA-ES is seeks to make the algorithm performance robust and not dependent on the objective function or tuning. The various *control parameters* of the algorithm were pre-defined on basis of theoretical aspects and practical benchmarks.

7 The Cartographer Class

Cartographer Class:

- I Use of Information: *implicit/explicit memory* by surrogate and solution archive
- II Candidate Evaluation: *indirect*
- III Type of Candidate: *single*
- IV Search Space: *global*
- V Problem Properties: multimodal, black-box, noisy, expensive

Cartographer algorithms differ from all other defined classes in their focus on acquiring, gathering and utilizing global information about the fitness landscape. They utilize prior evaluated candidates and model the acquired information to predict the fitness of new candidates. These models are then used for an efficient indirect search in which typically a *single* new candidate is proposed in each iteration, instead of performing multiple, direct and localized search steps.

Analogy 3 (The Cartographer).

The intuitive idea of the cartographer is a specialist who systematically measures a landscape by taking samples of the height to create a topological map. This map resembles the real landscape with a given approximation accuracy and is typically exact at the sampled locations (if the measurements are without variance) and models the remaining landscape by regression. It can then be examined and utilized by any other individual, such as a wanderer or guide, to find a desired location. One could think of a guide using a paper map or navigation system to find the place of interest.

As illustrated in Section 2.4, the surrogates $s(\mathbf{x})$ depict the *maps* of the fitness landscape of an objective function ($f_1(\mathbf{x})$ or $f_2(\mathbf{x})$) in an algorithmic framework. In this Section, we will first give a brief introduction to common surrogate models and then outline typical cartographer frameworks and their search strategies.

7.1 Surrogate Models

The surrogate is the core element of any surrogate-based optimization and essential for their performance. A perfect surrogate provides an excellent fit to prior observations, whilst ideally possessing superior interpolation and extrapolation abilities. However, the large number of available surrogate models all have significant differing characteristics, advantages and disadvantages. Model selection is thus a complicated and difficult task. If no domain knowledge is available, such as in real black box optimization, it is often inevitable to test different surrogates for their applicability.

Surrogates are built on basis of prior observations, which provide information about the fitness landscape of the problem. Thus, the initial candidates are commonly selected following different information criteria and suitable experimental design. For example, linear regression models can be build with factorial designs, while Gaussian process models are best coupled with space-filling designs, such as *Latin hypercube sampling* (Montgomery, Montgomery, and Montgomery, 1984; Sacks et al., 1989).

Common models are: linear, quadratic or polynomial regression, Gaussian processes (also known as Kriging) (Sacks et al., 1989; Forrester, Sobester, and Keane, 2008), regression trees (Breiman, Friedman, Stone, and Olshen, 1984), artificial neural networks and radial basis function networks (Haykin, 2004; Hornik, Stinchcombe, and White, 1989) including deep learning networks (Collobert and Weston, 2008; Hinton, Deng, Yu, Dahl, Mohamed, Jaitly, Senior, Vanhoucke, Nguyen, Sainath, et al., 2012; Hinton, Osindero, and Teh, 2006) and symbolic regression models (Augusto and Barbosa, 2000; Flasch, Mersmann, and Bartz-Beielstein, 2010; McKay, Willis, and Barton, 1995), which are usually optimized by genetic programming (Koza, 1992).

Further, a lot of effort in current studies is to research the benefits of model ensembles, which combine several distinct models (Goel, Haftka, Shyy, and Queipo, 2007; Müller and Shoemaker, 2014; Friese, Bartz-Beielstein, and Emmerich, 2016). The goal is to create a sophisticated predictor that surpasses the performance of a single model. A well-known example are random forests (Freund and Schapire, 1997), which use *bagging* to fit a large number of decision trees (Breiman, 2001). We regard ensemble modeling as state-of-the-art of current research, as they are able to combine the advantages of different models to generate outstanding results in both classification and regression. The drawback of these ensemble methodologies is that they are computational expensive and pose a demanding problem in regard to efficient model selection, evaluation and combination.

7.2 Cartographer Algorithms

Cartographer algorithms are surrogate-based optimization methodologies, which explicitly use a surrogate in their optimization cycle, following the general principle outlined in Section 2.4. They are either fixed algorithms designed around a certain model, such as *Kriging* in the well-known *efficient global optimization* (EGO) (Jones et al., 1998), or frameworks with a choice of

possible surrogates and optimization methods. We present two common frameworks and discuss their particularities: general *bayesian optimization* (Mockus, 1974) and *sequential parameter optimization* (Bartz-Beielstein et al., 2005; Bartz-Beielstein, 2010).

Forrester and Keane (2009) and Bartz-Beielstein and Zaefferer (2017) give overviews of surrogate-based optimization, different surrogate models and infill criteria. Moreover, they match surrogates to problem classes and give hints about their individual applicability. In general, the selection of an adequate model, experimental design and optimizer requires both domain knowledge and expertise. We will focus on the above-mentioned frameworks as they deliver a good, yet not complete, view of the surrogate-based search strategy.

7.2.1 Efficient Global Optimization

EGO (Jones et al., 1998) was motivated by the urge to develop a methodology to optimize expensive black-box functions. It utilizes Kriging surrogates and motivates the use *expected improvement* as infill criteria. In general, the algorithm consists of two phases: first, the initialization by Latin hypercube sampling and the construction of a Kriging surrogate; second, the iterative improvement of the best solution utilizing the surrogate. Before starting the second phase, Jones et al. (1998) suggests to analyze the surrogate model fit. If the fit is not satisfactory, it can be tried to improve it by a tuning of the model parameters or transformation of the data.

Algorithm 7.1: Efficient Global Optimization

```

1 begin
   | // phase 1: initial surrogate building
2   |  $t = 1$ 
3   | initialize  $k$  candidate solutions  $\{\mathbf{x}_{1:k}\}$  by latin hypercube sampling
4   | evaluate them on objective function  $\mathbf{y}_i = f_1(\mathbf{x}_i)$  or  $\mathbf{y}_i = f_2(\mathbf{x}_i), 1 \leq i \leq k$ ;
5   | build initial Kriging surrogate  $s_t$  with initial observations  $\mathcal{D}_t = \{(\mathbf{x}_i, \mathbf{y}_i), 1 \leq i \leq k\}$ 
6   | analyze and improve model fit (optional)
   | // phase 2: use and update surrogate
7   | while not termination-condition do
8     | if  $t > 1$  then
9       |   update the Kriging surrogate  $s_t$  with the set of all observations  $\mathcal{D}_t$ 
10      | end
11     | calculate expected improvement infill criteria on surrogate  $s_t$ 
12     | optimize EI for maximum by branch and bound; use optimum as candidate  $\mathbf{x}_t$ 
13     | evaluate  $\mathbf{x}_t$  on the objective function  $\mathbf{y}_t = f_1(\mathbf{x}_t)$  or  $\mathbf{y}_t = f_2(\mathbf{x}_t)$ 
14     | add new solutions to the set of all observations  $\mathcal{D}_{t+1} = \{\mathcal{D}_t, (\mathbf{x}_t, \mathbf{y}_t)\}$ 
15     |  $t = t + 1$ 
16   | end
17 end

```

The second phase starts the iterative optimization process, as described in Section 2.4. During the *variation*, a new candidate is searched by optimizing the expected improvement infill criteria of the Kriging surrogate. For this optimization, the exact *branch and bound* method is used (cf. exact class, Section 4). Expected improvement is motivated as infill criteria, because it guarantees a balance of exploration and exploitation by utilizing both the predicted best mean

value of the model, as well as the model uncertainty. An example for the complete methodology is outlined in Algorithm 7.1. The search strategy of EGO is a fundamental example for most surrogate-based optimization which are applicable to expensive optimization problems. However, we suggest to use advanced frameworks based on this base version of EGO. These frameworks are more flexible and applicable to a larger class of problem.

7.2.2 Bayesian Optimization

The term Bayesian optimization (BO) was introduced by Mockus (1974, 1994, 2012) and describes not a single, but a scheme for algorithms, which we regard as surrogate-based, particular these based on Gaussian processes. While the general BO scheme thus remains similar to the outlined algorithm in Section 2.4, BO differs in the underlying terminology: In BO, the user selects and initial, so-called *prior* distribution, which should support the *a priori* beliefs about the underlying unknown objective function. Gaussian distributions are suggested and a common choice. Algorithm 7.2 displays a general BO algorithm.

Algorithm 7.2: Bayesian Optimization

```

1 begin
  // phase 1: initialization
2    $t = 1$ 
3   select prior distribution  $s_{t-1}(\mathbf{x})$ 
4   initialize  $k$  candidate solutions  $\{\mathbf{x}_{1:k}\}$ 
5   evaluate them on objective function  $\mathbf{y}_i = f_1(\mathbf{x}_i)$  or  $\mathbf{y}_i = f_2(\mathbf{x}_i), 1 \leq i \leq k$ 
6   update prior with  $\mathcal{D}_t = \{(\mathbf{x}_i, \mathbf{y}_i), 1 \leq i \leq k\}$  to get initial posterior distribution  $s_t(\mathbf{x})$ ;
  // phase 2: optimization
7   while not termination-condition do
8     if  $t > 1$  then
9       | update posterior distribution  $s_t$  with the set of all observations  $\mathcal{D}_t$ 
10      end
11     calculate infill function on posterior distribution  $s_t(\mathbf{x})$ 
12     optimize infill criterion; use optimum as candidate  $\mathbf{x}_t$ 
13     evaluate  $\mathbf{x}_t$  on the objective function  $\mathbf{y}_t = f_1(\mathbf{x}_t)$  or  $\mathbf{y}_t = f_2(\mathbf{x}_t)$ 
14     add new solutions to the set of all observations  $\mathcal{D}_{t+1} = \{\mathcal{D}_t, (\mathbf{x}_t, \mathbf{y}_t)\}$ 
15      $t = t + 1$ 
16   end
17 end

```

This prior distribution is updated by sampled observations to acquire the *posterior* distribution. The optimization cycle includes the optimization of the *acquisition* (or infill) function to maximize *utility* or minimize *risk* (fitness). Typical choices include the probability of improvement (Kushner, 1964), expected improvement (Jones et al., 1998) and confidence bounds (Cox and John, 1997). Algorithms such as EGO can be seen as applied variants of BO. It is widely applicable to different applications, including expensive optimization problems (Lizotte, 2008; Khan, Goldberg, and Pelikan, 2002) and machine learning (Snoek, Larochelle, and Adams, 2012; Swersky, Snoek, and Adams, 2013). Brochu, Cora, and De Freitas (2010) give a tutorial on BO with different application examples.

7.2.3 Sequential Parameter Optimization Toolbox

The *sequential parameter optimization toolbox* (SPOT), developed by Bartz-Beielstein (2010), is a dynamic surrogate-based optimization framework, which was initially intended for *offline tuning* of algorithm *control parameters*. Various methods for initial sampling designs, different models and optimization techniques are included.

SPOT is strongly influenced by statistical methods from design of experiments, where it is attempted to prove a certain statistical hypothesis on basis of testing. Hereby, the available budget for experiments (i.e., function evaluations) is used sequentially to improve a solution and update the surrogate. This is done until sufficient knowledge about the search space is available to accept or reject the initial stated hypothesis.

The overall design was dedicated to algorithm tuning and follows two goals: one was improving the *efficiency* of an algorithm, i.e., discovering the algorithm parameters to solve a defined problem instance as fast as possible. The other goal was improving the *robustness* of an identified parameter setup, i.e., for solving different problem instances which for example differ in their region of interest or search space dimensionality. To tune stochastic algorithms, SPOT integrates *noise handling* techniques by dynamic re-sampling of solutions. This design can be transferred to general surrogate-based optimization, as the methods tackle the present balancing problem of exploration (*robustness*) and exploitation (*efficiency*).

The general framework of SPOT is similar to the general surrogate-based optimization algorithms EGO and BO, which are divided in initialization and iterative optimization phases. SPOT also explicitly includes a prior parametrization phase, where the user has to choose the surrogate and initial sampling design. SPOT defines a flexible framework and is thus applicable to a large range of problems, such as the mentioned algorithm tuning and industrial optimization.

8 The Hybrid Class

Hybrid Class:

- I Use of Information: *implicit/explicit memory* by surrogate or candidate archive
- II Candidate Evaluation: *direct/indirect*
- III Type of Candidate: *single ,population, distribution model*
- IV Search Space: *global*
- V Problem Properties: multimodal, black-box, noisy, discontinuities, multi-objective, expensive

The hybrid class depicts combinations of algorithms from the previously mentioned classes. Hybrid algorithms were developed as a strategy to improve or tackle individual algorithm weaknesses. The algorithms are often given distinctive roles of exploration and exploitation, as they are combinations of an explorative global search method paired with a local search algorithm. For example, population-based algorithms with remarkable exploration abilities can be paired with local algorithms with fast convergence. This approach gives some benefits, as the combined algorithms can be adapted or tuned to fulfill their distinct tasks. Moreover, the concepts can be easily adapted for parallel frameworks.

One of the most successful type of hybrids are the *surrogate-assisted evolutionary algorithms* (Emmerich, Giannakoglou, and Naujoks, 2006; Lim, Jin, Ong, and Sendhoff, 2010). An overview

of surrogate-assisted optimization is given by Jin (2011), including several examples for real-world applications. An example is outlined in Algorithm 8.1.

Algorithm 8.1: Surrogate-Assisted Evolutionary Algorithm

```

1 begin
2   t=0
3   initialize  $\mu$  random population  $\mathcal{P}_t = \{\mathbf{x}_{t,1:\mu}\} \subseteq \mathcal{S}$ 
4   evaluate population  $\mathcal{P}_t$ 
5   while not termination-condition do
6     run evolutionary algorithm to find  $\lambda$  offspring  $\mathcal{O}_t = \{\mathbf{x}_i, \mu < i \leq \mu + \lambda\}$ 
7     build surrogate  $s_t(\mathbf{x})$  ( with current observations  $\mathcal{D}_t = \{(\mathbf{x}_i, \mathbf{y}_i), 1 \leq i \leq k\}$ )
8     predict fitness of offspring  $\mathcal{O}_t$  using surrogate  $s_t(\mathbf{x})$ 
9     select  $\lambda^*$  best offspring  $\mathcal{O}_t^*$  based on predicted fitness
10    evaluate selected offspring  $\mathcal{O}_t^*$  with objective function
11    run local optimizer from  $\mathbf{x} \in \mathcal{O}_t^*$  as starting solution to get refined solutions  $\mathcal{O}_t^{**}$ 
12    select survivors from  $\mathcal{P}_t \cup \mathcal{O}_t^* \cup \mathcal{O}_t^{**}$  for next generation  $\mathcal{P}_{t+1}$ 
13    t = t + 1
14  end
15 end

```

Example 8.1 (Surrogate-Assisted Evolutionary Algorithm). In this hybrid search strategy, a local surrogate is built upon the current parent population and utilized to predict the fitness of a number of λ offspring candidates. The selection is then based on the predicted fitness of the surrogate. Optionally, a local optimizer can be used to further refine the computed solutions. Extensive use of local search leads to a fast convergence to local optima. This hybrid strategy can be altered by using the surrogate only for a part of the generated offspring, while the other part is evaluated with the real fitness function.

Additional examples of hybrids can be found in the literature, covering all possible class combinations:

Memetic algorithms, as defined by Moscato et al. (1989), are a class of search methods which combine population-based guides with a local wanderer. An extensive overview of memetic algorithms is given by Molina, Lozano, García-Martínez, and Herrera (2010). They describe how different hybrid algorithms can be constructed by looking at suitable local search algorithms with special regard to their convergence abilities.

Bartz-Beielstein, Preuss, and Rudolph (2006) describe a hybrid approach an evolution strategy for exploration with quasi-Newton method for exploitation. The algorithm runs the ES and the local search in a consecutive way and the budget of evaluations for each method is controlled by a control parameter. They performed experiments in which they varied the budget parameter to test if this hybrid approach can be superior to running both methods individually. They came to the conclusion, that hybridization can be beneficial for difficult objective functions, as the ES provides information about interesting regions where the local search is then applied.

The *surrogate management framework* (Booker et al., 1999; Serafini, 1999) utilizes a combination of a global surrogate-based algorithm with an exact local pattern search. Moreover, it uses fitness space transformations from continuous to combinatorial by introducing a finite mesh

of possible solutions. The key concept is to lower the optimization costs by reducing the real function evaluations using a surrogate, while at the same time retaining the benefits of a combinatorial search space and pattern search, i.e., the robust convergence behavior. The algorithm utilizes two distinct global and local search phases, which are executed during the sequential optimization. In the global *search* step the selected infill criterion (e.g. expected improvement) is optimized in continuous space and the nearest mesh point is selected. In the local *poll* step, a set of candidates situated on the mesh around the current best solution is evaluated directly on the objective.

Taddy, Lee, Gray, and Griffin (2009) combine surrogate-based optimization based on *treed Gaussian processes* (TGP) with exact *asynchronous parallel pattern search* (APPS) in a parallel search framework. The algorithm starts with a space filling initial sampling using a latin hypercube design, then runs TGP and APPS in parallel. In this case, TGP is used to predict a ranked list of a fixed number of new candidates, while APPS is used to perform local optimization runs. The budget for evaluation and computation is split between these two components and all observations are stored in an shared archive.

Hybrid algorithms are applicable to a large class of problems, defined by which class their algorithms originate from. Their downside is their large complexity and the risk, that their higher complexity does not lead to improved performance, due to the difficult balancing and required tuning of the distinct algorithms. Their complex search strategies with a large number of control parameter can make them difficult to tune. The algorithm itself becomes a black-box, as the underlying search strategy and the convergence behavior is influenced by numerous operators and difficult to comprehend.

9 Concluding Remarks

In this work, we presented an overview of continuous global optimization algorithms with focus on explaining their search strategies using a new intuitive taxonomy. We defined a set of five classes: *exact*, *wanderer*, *guide*, *cartographer*, and *hybrid* in the Sections 4 to 8 and outlined their individual properties and exemplified algorithms for each of the proposed classes.

To shortly recapitulate: the *exact* class utilizes (a priori) problem information to solve a problem with a guarantee of finding the optimum. The heuristic search strategies of the class of *wanderer* algorithms are suitable for fast convergence in a unimodal search space and often part of other algorithms. The well-known metaheuristics from the *guide* class were developed for general applicability, particular for multi-modal problems with unknown properties. The *cartographer* class focuses on surrogate-based algorithms and regarded frameworks for problems with expensive function evaluations. Last but not least, we took a look at *hybrid* class algorithms, which try to combine the strengths of different algorithms to overcome their individual weaknesses.

In general, it is beneficial for each user to identify if an optimization algorithm is suitable for their problem before applying them. To support users in selecting a suitable algorithm, we pointed out the pros and cons of the different search strategies, the individual algorithm features and typical characteristics of CGO problems they are able to handle.

At this point, we also want to highlight a new promising research area in the field of algorithms, which is *automated algorithm selection* and particular *automatic algorithm configuration*. Both ideas tackle the problem of selecting the correct search strategy for a given problem. Automated algorithm selection tries to find the most suitable algorithm for a certain problem based on

machine learning and exploited problem information, such as explorative landscape analysis. This method of algorithm selection has shown to be able to outperform a single algorithm on a set of benchmark functions (Kerschke and Trautmann, 2017). An even more promising result was presented by van Rijn, Wang, van Stein, and Bäck (2017), where *algorithm configuration* was used to select algorithmic components for creating a search strategy outperforming available algorithms. The particular interesting idea is here, that search operators of algorithms are identified, extracted and then again combined to a new search strategy. The whole procedure also shows the strong connections between different named algorithms, particular in the area of bio-inspired metaheuristics.

Interesting challenges for future algorithm design arise from problems in engineering applications, where the data available is restricted to certain conditions, such as streaming and online data and dynamic problems. The need for new optimization approaches emerges from rapid development of communicating sensors and machines in the field of engineering, also known as *internet of things* (Atzori, Iera, and Morabito, 2010). Suitable optimization algorithms need to be directly included in the production cycle, adapting to generate robust solutions in challenging dynamic environments with moving optima. Dynamic, surrogate-based online learning, where a complex static surrogate is constructed and combined with time-varying modeling, is still an open issue (Jin and Branke, 2005). Nowadays, cloud computing and high-performance computing clusters are available to a wide range of users. Many optimization algorithms are not yet fitted for the needs of parallel computation and need to be adapted Rehbach, Zaefferer, Stork, and Bartz-Beielstein (2018). The large and successful field of deep learning networks (LeCun, Bengio, and Hinton, 2015; Schmidhuber, 2015) declares a complete new field from which very complex and difficult optimization problems arise. The extension of surrogate-based optimization to these fields, e.g., parallel frameworks and deep learning, is an interesting research topic. Further, we identify a lack in the field of realistic benchmarks, which are based on real-world data sets, which would allow a realistic comparison of different algorithmic approaches.

10 Acknowledgements

Thanks to all our colleagues for helpful proofreading, remarks and writing suggestions. This work is part of a project that has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement no. 692286.

References

- F. Archetti and F. Schoen. A survey on the global optimization problem: general theory and computational approaches. *Annals of Operations Research*, 1(2):87–110, 1984.
- D. V. Arnold and H.-G. Beyer. A comparison of evolution strategies with other direct search methods in the presence of noise. *Computational Optimization and Applications*, 24(1):135–159, 2003.
- D. V. Arnold and N. Hansen. A (1+ 1)-cma-es for constrained optimisation. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, pages 297–304. ACM, 2012.
- J. Arora, O. Elwakeil, A. Chahande, and C. Hsieh. Global optimization methods for engineering applications: a review. *Structural optimization*, 9(3-4):137–159, 1995.

- L. Atzori, A. Iera, and G. Morabito. The internet of things: A survey. *Computer networks*, 54(15): 2787–2805, 2010.
- C. Audet. A survey on direct search methods for blackbox optimization and their applications. In *Mathematics Without Boundaries*, pages 31–56. Springer, 2014.
- D. A. Augusto and H. J. Barbosa. Symbolic regression via genetic programming. In *Neural Networks, 2000. Proceedings. Sixth Brazilian Symposium on*, pages 173–178. IEEE, 2000.
- T. Back. *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms*. Oxford university press, 1996.
- T. Bartz-Beielstein. Spot: An r package for automatic and interactive tuning of optimization algorithms by sequential parameter optimization. *arXiv preprint arXiv:1006.4645*, 2010.
- T. Bartz-Beielstein and M. Zaefferer. Model-based methods for continuous and discrete global optimization. *Applied Soft Computing*, 55:154–167, 2017.
- T. Bartz-Beielstein, C. W. Lasarczyk, and M. Preuß. Sequential parameter optimization. In *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, volume 1, pages 773–780. IEEE, 2005.
- T. Bartz-Beielstein, M. Preuss, and G. Rudolph. Investigation of one-go evolution strategy/quasi-newton hybridizations. In *Hybrid Metaheuristics*, pages 178–191. Springer, 2006.
- H.-G. Beyer. *The theory of evolution strategies*. Springer Science & Business Media, 2013.
- A. J. Booker, J. Dennis Jr, P. D. Frank, D. B. Serafini, V. Torczon, and M. W. Trosset. A rigorous framework for optimization of expensive functions by surrogates. *Structural optimization*, 17(1):1–13, 1999.
- I. Boussaïd, J. Lepagnot, and P. Siarry. A survey on optimization metaheuristics. *Information Sciences*, 237:82–117, 2013.
- L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, Oct 2001. ISSN 1573-0565. doi: 10.1023/A:1010933404324. URL <https://doi.org/10.1023/A:1010933404324>.
- L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen. *Classification and regression trees*. CRC press, 1984.
- E. Brochu, V. M. Cora, and N. De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
- Y. Carson and A. Maria. Simulation optimization: methods and applications. In *Proceedings of the 29th conference on Winter simulation*, pages 118–126. IEEE Computer Society, 1997.
- R. Chelouah and P. Siarry. Tabu search applied to global optimization. *European journal of operational research*, 123(2):256–270, 2000.

- C. A. C. Coello. Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art. *Computer methods in applied mechanics and engineering*, 191(11):1245–1287, 2002.
- R. Collobert and J. Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167. ACM, 2008.
- A. Corana, M. Marchesi, C. Martini, and S. Ridella. Minimizing multimodal functions of continuous variables with the “simulated annealing” algorithm. *ACM Transactions on Mathematical Software (TOMS)*, 13(3):262–280, 1987.
- D. D. Cox and S. John. Sdo: A statistical method for global optimization. *Multidisciplinary design optimization: state of the art*, pages 315–329, 1997.
- M. Črepinšek, S.-H. Liu, and M. Mernik. Exploration and exploitation in evolutionary algorithms: a survey. *ACM Computing Surveys (CSUR)*, 45(3):35, 2013.
- M. Dorigo, M. Birattari, and T. Stutzle. Ant colony optimization. *IEEE computational intelligence magazine*, 1(4):28–39, 2006.
- A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer, 2nd edition, 2015. ISBN 3662448734, 9783662448731.
- A. E. Eiben, R. Hinterding, and Z. Michalewicz. Parameter control in evolutionary algorithms. *Evolutionary Computation, IEEE Transactions on*, 3(2):124–141, 1999.
- M. T. Emmerich, K. C. Giannakoglou, and B. Naujoks. Single- and multiobjective evolutionary optimization assisted by gaussian random field metamodelling. *IEEE Transactions on Evolutionary Computation*, 10(4):421–439, 2006.
- R. Faber, T. Jockenhövel, and G. Tsatsaronis. Dynamic optimization with simulated annealing. *Computers & chemical engineering*, 29(2):273–290, 2005.
- P. Filipiak and P. Lipinski. Infeasibility driven evolutionary algorithm with feed-forward prediction strategy for dynamic constrained optimization problems. In *Applications of Evolutionary Computation*, pages 817–828. Springer, 2014.
- O. Flasch, O. Mersmann, and T. Bartz-Beielstein. Rgp: An open source genetic programming system for the r environment. In *Proceedings of the 12th Annual Conference Companion on Genetic and Evolutionary Computation, GECCO '10*, pages 2071–2072, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0073-5. doi: 10.1145/1830761.1830867. URL <http://doi.acm.org/10.1145/1830761.1830867>.
- P. J. Fleming and R. C. Purshouse. Evolutionary algorithms in control systems engineering: a survey. *Control engineering practice*, 10(11):1223–1241, 2002.
- C. A. Floudas. *Deterministic global optimization: theory, methods and applications*, volume 37. Springer Science & Business Media, 2013.
- L. J. Fogel, A. J. Owens, and M. J. Walsh. *Artificial intelligence through simulated evolution*. John Wiley, 1966.

- F. V. Fomin and P. Kaski. Exact exponential algorithms. *Commun. ACM*, 56(3):80–88, Mar. 2013. ISSN 0001-0782. doi: 10.1145/2428556.2428575. URL <http://doi.acm.org/10.1145/2428556.2428575>.
- C. M. Fonseca, P. J. Fleming, et al. Genetic algorithms for multiobjective optimization: Formulation discussion and generalization. In *Icga*, volume 93, pages 416–423, 1993.
- A. Forrester, A. Sobester, and A. Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons, Inc., 2008.
- A. I. Forrester and A. J. Keane. Recent advances in surrogate-based optimization. *Progress in Aerospace Sciences*, 45(1):50–79, 2009.
- Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- M. Friese, T. Bartz-Beielstein, and M. T. M. Emmerich. Building ensembles of surrogates by optimal convex combination. In M. Mernik and G. Papa, editors, *Bioinspired Optimization Methods and their Applications*, pages 131–144, May 2016.
- R. Gämperle, S. D. Müller, and P. Koumoutsakos. A parameter study for differential evolution. *Advances in intelligent systems, fuzzy systems, evolutionary computation*, 10:293–298, 2002.
- F. Glover. Tabu search—part i. *ORSA Journal on computing*, 1(3):190–206, 1989.
- T. Goel, R. T. Haftka, W. Shyy, and N. V. Queipo. Ensemble of surrogates. *Structural and Multidisciplinary Optimization*, 33(3):199–216, 2007.
- W. L. Goffe, G. D. Ferrier, and J. Rogers. Global optimization of statistical functions with simulated annealing. *Journal of Econometrics*, 60(1):65–99, 1994.
- E. Hansen and G. W. Walster. *Global optimization using interval analysis: revised and expanded*, volume 264. CRC Press, 2003.
- N. Hansen, S. D. Müller, and P. Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es). *Evolutionary computation*, 11(1):1–18, 2003.
- N. Hansen, A. Auger, R. Ros, S. Finck, and P. Pošík. Comparing results of 31 algorithms from the black-box optimization benchmarking bbob-2009. In *Proceedings of the 12th annual conference companion on Genetic and evolutionary computation*, pages 1689–1696. ACM, 2010a.
- P. Hansen and N. Mladenovic. Variable neighbourhood search. *Handbook of Metaheuristics*, Dordrecht, Kluwer Academic Publishers, 2003.
- P. Hansen, N. Mladenović, and J. A. M. Pérez. Variable neighbourhood search: methods and applications. *Annals of Operations Research*, 175(1):367–407, 2010b.
- M. Hauschild and M. Pelikan. An introduction and survey of estimation of distribution algorithms. *Swarm and Evolutionary Computation*, 1(3):111–128, 2011.
- S. Haykin. A comprehensive foundation. *Neural Networks*, 2(2004), 2004.

- D. Henderson, S. H. Jacobson, and A. W. Johnson. The theory and practice of simulated annealing. In *Handbook of metaheuristics*, pages 287–319. Springer, 2003.
- G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6): 82–97, 2012.
- G. E. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.
- J. H. Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. 1992.
- K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- R. Horst and H. Tuy. *Global optimization: Deterministic approaches*. Springer Science & Business Media, 2013.
- N. Hu. Tabu search method with random moves for globally optimal design. *International Journal for Numerical Methods in Engineering*, 35(5):1055–1070, 1992.
- F. Hutter, H. H. Hoos, and K. Leyton-Brown. Sequential model-based optimization for general algorithm configuration. *LION*, 5:507–523, 2011.
- Y. Jin. Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation*, 1(2):61–70, 2011.
- Y. Jin and J. Branke. Evolutionary optimization in uncertain environments—a survey. *IEEE Transactions on evolutionary computation*, 9(3):303–317, 2005.
- D. R. Jones. A taxonomy of global optimization methods based on response surfaces. *Journal of global optimization*, 21(4):345–383, 2001.
- D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the lipschitz constant. *Journal of optimization Theory and Applications*, 79(1):157–181, 1993.
- D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.
- J. Kennedy and R. Eberhart. Particle swarm optimization. In *Neural Networks, 1995. Proceedings., IEEE International Conference on*, volume 4, pages 1942–1948. IEEE, 1995.
- P. Kerschke and H. Trautmann. Automated algorithm selection on continuous black-box problems by combining exploratory landscape analysis and machine learning. *arXiv preprint arXiv:1711.08921*, 2017.
- N. Khan, D. E. Goldberg, and M. Pelikan. Multi-objective bayesian optimization algorithm. In *Proceedings of the 4th Annual Conference on Genetic and Evolutionary Computation, GECCO’02*, pages 684–684, San Francisco, CA, USA, 2002. Morgan Kaufmann Publishers Inc. ISBN 1-55860-878-8. URL <http://dl.acm.org/citation.cfm?id=2955491.2955599>.

- S. Kirkpatrick, C. D. Gelatt, M. P. Vecchi, et al. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.
- T. G. Kolda, R. M. Lewis, and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM review*, 45(3):385–482, 2003.
- J. R. Koza. *Genetic programming: on the programming of computers by means of natural selection*, volume 1. MIT press, 1992.
- H. J. Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 86(1):97–106, 1964.
- P. Larrañaga and J. A. Lozano. *Estimation of distribution algorithms: A new tool for evolutionary computation*, volume 2. Springer Science & Business Media, 2001.
- E. L. Lawler and D. E. Wood. Branch-and-bound methods: A survey. *Operations research*, 14(4): 699–719, 1966.
- Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- A. Leon. A classified bibliography on optimization. *Recent Advances in Optimization Techniques*, John Wiley & Sons, Inc., New York and London, pages 599–649, 1966.
- R. M. Lewis, V. Torczon, and M. W. Trosset. Direct search methods: then and now. *Journal of computational and Applied Mathematics*, 124(1):191–207, 2000.
- D. Lim, Y. Jin, Y.-S. Ong, and B. Sendhoff. Generalizing surrogate-assisted evolutionary computation. *IEEE Transactions on Evolutionary Computation*, 14(3):329–355, 2010.
- D. J. Lizotte. *Practical bayesian optimization*. University of Alberta, 2008.
- M. Locatelli. Simulated annealing algorithms for continuous global optimization. In *Handbook of global optimization*, pages 179–229. Springer, 2002.
- M. López-Ibañez, J. Dubois-Lacoste, L. P. Cáceres, M. Birattari, and T. Stützle. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3:43–58, 2016.
- I. Loshchilov, M. Schoenauer, and M. Sebag. Self-adaptive surrogate-assisted covariance matrix adaptation evolution strategy. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, pages 321–328. ACM, 2012.
- R. T. Marler and J. S. Arora. Survey of multi-objective optimization methods for engineering. *Structural and multidisciplinary optimization*, 26(6):369–395, 2004.
- A. L. Marsden, M. Wang, J. E. Dennis Jr, and P. Moin. Optimal aeroacoustic shape design using the surrogate management framework. *Optimization and Engineering*, 5(2):235–262, 2004.
- B. McKay, M. J. Willis, and G. W. Barton. Using a tree structured genetic algorithm to perform symbolic regression. In *Genetic Algorithms in Engineering Systems: Innovations and Applications*, 1995. GALEZIA. First International Conference on (Conf. Publ. No. 414), pages 487–492. IET, 1995.

- Z. Michalewicz and D. B. Fogel. *How to solve it: modern heuristics*. Springer Science & Business Media, 2013.
- N. Mladenović, M. Dražić, V. Kovačević-Vujčić, and M. Čangalović. General variable neighborhood search for the continuous optimization. *European Journal of Operational Research*, 191(3): 753–770, 2008.
- J. Mockus. On bayesian methods for seeking the extremum. In *Proceedings of the IFIP Technical Conference*, pages 400–404. Springer-Verlag, 1974.
- J. Mockus. Application of bayesian approach to numerical methods of global and stochastic optimization. *Journal of Global Optimization*, 4(4):347–365, 1994.
- J. Mockus. *Bayesian approach to global optimization: theory and applications*, volume 37. Springer Science & Business Media, 2012.
- D. Molina, M. Lozano, C. García-Martínez, and F. Herrera. Memetic algorithms for continuous optimisation based on local search chains. *Evolutionary Computation*, 18(1):27–63, 2010.
- D. C. Montgomery, D. C. Montgomery, and D. C. Montgomery. *Design and analysis of experiments*, volume 7. Wiley New York, 1984.
- P. Moscato et al. On evolution, search, optimization, genetic algorithms and martial arts: Towards memetic algorithms. *Caltech concurrent computation program, C3P Report*, 826:1989, 1989.
- J. Müller and C. A. Shoemaker. Influence of ensemble surrogate models and sampling strategy on the solution quality of algorithms for computationally expensive black-box global optimization problems. *Journal of Global Optimization*, 60(2):123–144, 2014.
- B. Naujoks, N. Beume, and M. Emmerich. Multi-objective optimisation using s-metric selection: Application to three-dimensional solution spaces. In *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, volume 2, pages 1282–1289. IEEE, 2005.
- A. Neumaier. Complete search in continuous global optimization and constraint satisfaction. *Acta numerica*, 13:271–369, 2004.
- Y. S. Ong, P. Nair, A. Keane, and K. Wong. Surrogate-assisted evolutionary optimization frameworks for high-fidelity engineering design problems. In *Knowledge Incorporation in Evolutionary Computation*, pages 307–331. Springer, 2005.
- J. Pearl. Heuristics. intelligent search strategies for computer problem solving. *The Addison-Wesley Series in Artificial Intelligence, Reading, Mass.: Addison-Wesley, 1985, Reprinted version*, 1985.
- J. D. Pintér. Global optimization: software, test problems, and applications. In *Handbook of global optimization*, pages 515–569. Springer, 2002.
- N. V. Queipo, R. T. Haftka, W. Shyy, T. Goel, R. Vaidyanathan, and P. K. Tucker. Surrogate-based analysis and optimization. *Progress in aerospace sciences*, 41(1):1–28, 2005.

- I. Rechenberg. *Evolutionsstrategie–Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboog, 1973.
- I. Rechenberg. *Evolutionsstrategie '94*. Frommann-Holzboog, 1994.
- F. Rehbach, M. Zaefferer, J. Stork, and T. Bartz-Beielstein. Comparison of parallel surrogate-assisted optimization approaches. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1348–1355. ACM, 2018.
- G. Rozenberg, T. Bck, and J. N. Kok. *Handbook of natural computing*. Springer Publishing Company, Incorporated, 2011.
- J. Sacks, W. J. Welch, T. J. Mitchell, and H. P. Wynn. Design and analysis of computer experiments. *Statistical science*, pages 409–423, 1989.
- J. Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- M. Schonlau. *Computer Experiments and Global Optimization*. PhD thesis, University of Waterloo, Ontario, Canada, 1997.
- H.-P. Schwefel. *Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie: mit einer vergleichenden Einführung in die Hill-Climbing-und Zufallsstrategie*. Birkhäuser, 1977.
- H.-P. P. Schwefel. *Evolution and optimum seeking: the sixth generation*. John Wiley & Sons, Inc., 1993.
- D. B. Serafini. *A framework for managing models in nonlinear optimization of computationally expensive functions*. PhD thesis, Rice University, 1999.
- D. F. Shanno. Conditioning of quasi-newton methods for function minimization. *Mathematics of computation*, 24(111):647–656, 1970.
- Y. Shi and R. Eberhart. A modified particle swarm optimizer. In *Evolutionary Computation Proceedings, 1998. IEEE World Congress on Computational Intelligence., The 1998 IEEE International Conference on*, pages 69–73. IEEE, 1998.
- O. M. Shir and T. Bäck. Niching in evolution strategies. In *Proceedings of the 7th annual conference on Genetic and evolutionary computation*, pages 915–916. ACM, 2005.
- P. Siarry and G. Berthiau. Fitting of tabu search to optimize functions of continuous variables. *International journal for numerical methods in engineering*, 40(13):2449–2457, 1997.
- P. Siarry, G. Berthiau, F. Durdin, and J. Haussy. Enhanced simulated annealing for globally minimizing functions of many-continuous variables. *ACM Transactions on Mathematical Software (TOMS)*, 23(2):209–228, 1997.
- S. Smit and A. Eiben. Multi-problem parameter tuning using bonesa. In *Artificial Evolution*, pages 222–233, 2011.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.

- K. Socha and M. Dorigo. Ant colony optimization for continuous domains. *European journal of operational research*, 185(3):1155–1173, 2008.
- J. Søndergaard, K. Madsen, and H. B. Nielsen. *Optimization using surrogate models-by the Space Mapping technique*. PhD thesis, Technical University of Denmark, 2003.
- K. Swersky, J. Snoek, and R. P. Adams. Multi-task bayesian optimization. In *Advances in neural information processing systems*, pages 2004–2012, 2013.
- M. A. Taddy, H. K. Lee, G. A. Gray, and J. D. Griffin. Bayesian guided pattern search for robust local optimization. *Technometrics*, 51(4):389–401, 2009.
- E.-G. Talbi. *Metaheuristics: from design to implementation*, volume 74. John Wiley & Sons, 2009.
- A. Törn and A. Zilinskas. *Global Optimization*. Springer, 1989.
- J. Van Groenigen and A. Stein. Constrained optimization of spatial sampling using continuous simulated annealing. *Journal of Environmental Quality*, 27(5):1078–1086, 1998.
- S. van Rijn, H. Wang, B. van Stein, and T. Bäck. Algorithm configuration data mining for cma evolution strategies. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '17*, pages 737–744, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4920-8. doi: 10.1145/3071178.3071205. URL <http://doi.acm.org/10.1145/3071178.3071205>.
- D. H. Wolpert and W. G. Macready. No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, Apr. 1997.
- K. S. Won and T. Ray. Performance of kriging and cokriging based surrogate models within the unified framework for surrogate assisted optimization. In *Evolutionary Computation, 2004. CEC2004. Congress on*, volume 2, pages 1577–1585. IEEE, 2004.
- S. H. Zanakis and J. R. Evans. Heuristic “Optimization”: Why, When, and How to Use It. *Interfaces*, 11(5):84–91, 1981. doi: 10.1287/inte.11.5.84. URL <http://dx.doi.org/10.1287/inte.11.5.84>.
- A. Žilinskas. A review of statistical models for global optimization. *Journal of Global Optimization*, 2(2):145–153, 1992.
- M. Zlochin, M. Birattari, N. Meuleau, and M. Dorigo. Model-based search for combinatorial optimization: A critical survey. *Annals of Operations Research*, 131(1-4):373–395, 2004.

Kontakt/Impressum

Diese Veröffentlichungen erscheinen im Rahmen der Schriftenreihe "Ciplus". Alle Veröffentlichungen dieser Reihe können unter

<https://cos.bibl.th-koeln.de/home>
abgerufen werden.

Die Verantwortung für den Inhalt dieser Veröffentlichung liegt beim Autor.

Datum der Veröffentlichung: 31.08.2018

Herausgeber / Editorship

Prof. Dr. Thomas Bartz-Beielstein,
Prof. Dr. Wolfgang Konen,
Prof. Dr. Boris Naujoks,
Prof. Dr. Horst Stenzel
Institute of Computer Science,
Faculty of Computer Science and Engineering Science,
TH Köln,
Steinmüllerallee 1,
51643 Gummersbach
url: www.ciplus-research.de

Schriftleitung und Ansprechpartner/ Contact editor's office

Prof. Dr. Thomas Bartz-Beielstein,
Institute of Computer Science,
Faculty of Computer Science and Engineering Science,
TH Köln,
Steinmüllerallee 1, 51643 Gummersbach
phone: +49 2261 8196 6391
url: <http://www.spotseven.de>
eMail: thomas.bartz-beielstein@th-koeln.de

ISSN (online) 2194-2870

**Technology
Arts Sciences
TH Köln**

