

# A Survey of Methods for Automated Algorithm Configuration

**Elias Schede**

*Decision and Operation Technologies Group,  
Bielefeld University, Bielefeld, Germany*

ELIAS.SCHEDE@UNI-BIELEFELD.DE

**Jasmin Brandt**

**Alexander Tornede**

*Department of Computer Science,  
Paderborn University, Paderborn, Germany*

JASMIN.BRANDT@UPB.DE

ALEXANDER.TORNEDE@UPB.DE

**Marcel Wever**

*Institute of Informatics, LMU Munich &  
Munich Center for Machine Learning, Munich, Germany*

MARCEL.WEVER@IFI.LMU.DE

**Viktor Bengs**

*Institute of Informatics,  
LMU Munich, Munich, Germany*

VIKTOR.BENGS@IFI.LMU.DE

**Eyke Hüllermeier**

*Institute of Informatics, LMU Munich &  
Munich Center for Machine Learning, Munich, Germany*

EYKE@LMU.DE

**Kevin Tierney**

*Decision and Operation Technologies Group,  
Bielefeld University, Bielefeld, Germany*

KEVIN.TIERNEY@UNI-BIELEFELD.DE

## Abstract

Algorithm configuration (AC) is concerned with the automated search of the most suitable parameter configuration of a parametrized algorithm. There is currently a wide variety of AC problem variants and methods proposed in the literature. Existing reviews do not take into account all derivatives of the AC problem, nor do they offer a complete classification scheme. To this end, we introduce taxonomies to describe the AC problem and features of configuration methods, respectively. We review existing AC literature within the lens of our taxonomies, outline relevant design choices of configuration approaches, contrast methods and problem variants against each other, and describe the state of AC in industry. Finally, our review provides researchers and practitioners with a look at future research directions in the field of AC.

## 1. Introduction

Difficult computational problems must be regularly solved in many areas of industry and academia, such as constraint satisfaction problems, Boolean satisfiability problems (SAT), vehicle routing problems, finding a proper machine learning model for a given dataset, or computing highly complex simulations. Algorithms that were developed to solve such problems usually have parameters that strongly influence the behavior of the respective algorithm and also, for example, the runtime that is required to solve problem instances

or the quality of returned solutions. In particular, different parameter values, also referred to as configurations in the following, are required for different sets of problem instances to achieve optimal results with respect to the running time or solution quality. It is, therefore, crucial to adapt the configuration of an algorithm to the given data or specifics of the set of problem instances in question. However, the configuration of algorithms, i.e., determining suitable parameter values, is a nontrivial and complex undertaking, since the algorithm must be actually executed for different configurations to observe the target metrics (e.g., runtime or solution quality).

The research field of algorithm configuration<sup>1</sup> (AC) has emerged in response to this problem. Especially within the last two decades, many approaches and problem variants have been proposed in this field. Generally speaking, the approaches try to find effective configurations of algorithms as efficiently as possible and recommend them for new, unseen problem instances. To this end, the approaches are typically given a training set of problem instances in an offline phase, which can be used as an input to run the algorithms, observe their performance, and (hopefully) generalize these observations to make good recommendations in production settings.

To illustrate the benefits of automatically configuring algorithm parameters, we provide the circuit satisfiability problem as an example. It refers to a classic SAT problem in which the task is to find an assignment of values such that the output of a Boolean circuit evaluates to true (Marques-Silva, 2008). In a business application, many such circuits must be evaluated to check their feasibility in limited time. This requires an efficient SAT solver such as Glucose (Audemard & Simon, 2009) to provide assignments in a timely manner. Glucose exposes several parameters that influence the search for assignments and, therefore, the time needed to evaluate a circuit.

Indeed, configurations of SAT solvers found by ParamILS (Hutter et al., 2007b), one of the first procedures to search for high-quality parameters in a structured way, achieve considerable speedups compared to default configurations of solvers. In particular, the configurations found by ParamILS reduce the arithmetic mean runtime for software verification instances for the SAT solver SPEAR from 787.1s to 1.5 seconds in the best case (Hutter et al., 2007a). More recently, PyDGGA (Ansótegui et al., 2021) reduced the solving time of the SAT solver SparrowToRiss (Balint & Manthey, 2013) on instances from the N-Rooks (Lindauer & Hutter, 2018a) dataset from 116 to 6.3 seconds. This shows that a practitioner will possibly achieve significant performance gains solving circuit assignments in the long run by only investing a limited amount of time into configuring Glucose for the specific task upfront. The recommendation of configurations for unseen problem instances can be seen as a common property of all algorithm configuration problem variants, which differ in terms of their concrete setting specifications. For example, there might only be a single (finite) categorical parameter to be configured, also referred to as algorithm selection (Rice, 1976), ranging from just a handful up to thousands of choices (Tornede et al., 2020b). In the other extreme, the space of algorithm configurations may take into account many parameters and could comprise infinitely many possible configurations. Beyond that, the problem variants differ in several other aspects, such as the objective function to be

---

1. In some works, the terms *parameter tuning* and *algorithm configuration* are used interchangeably (Hoos, 2011), however, we prefer *algorithm configuration*, following the argumentation of Hutter et al. (2009b) that this term implies a more general configuration setting.

optimized (runtime or solution quality), whether multiple objectives are to be considered simultaneously, whether training is performed offline or on the fly in an online setting, etc. Depending on the specific properties of an AC problem, algorithm configuration approaches, also referred to as algorithm configurators, can be more or less suited, or cannot be applied at all.

In this paper, we provide an overview of different variants of both AC problems and algorithm configurators. To this end, we propose two classification schemes: one for AC problems, and one for algorithm configurators. Based on this, we structure and summarize the available literature and classify existing problem variants as well as approaches to AC.

The remainder of the paper is structured as follows. First, in Section 2, we give a formal introduction into the setting of algorithm configuration, specify the scope of this survey, and discuss the relation between AC, AS and HPO. In Section 3, we present the classification schemes for AC problems and approaches that are used, in turn, to describe and compare existing algorithm configurators. In Sections 4 and 5, we survey algorithm configuration methods grouped by the property of whether these methods are model-free or leverage a model respectively. Section 6 deals with theoretical guarantees that can be obtained. Different problem variants, such as realtime AC, instance-specific vs. feature-based, multi-objective, and dynamic AC are discussed in Sections 7 to 10. Eventually, with the help of our classification schemes, we elaborate on appealing research directions in Section 11 and conclude this survey in Section 12. A list of abbreviations used in this work can be found in Table 6. In addition, we provide a list of useful software in Table 7. We note, however, that this list is by no means exhaustive; it is meant to provide an idea about available software at the time of publication.

## 2. Problem Formulation

### 2.1 Algorithm Configuration

To describe the AC problem more formally, we introduce the following notation that is similar to Hutter et al. (2009b). Let  $\mathcal{I}$  be a space of *problem instances* over which a *probability distribution*  $\mathcal{P}$  is defined. Optional *feature vectors*  $\mathbf{f}_i \in \mathbb{R}^d$  with features  $f_{i,1}, \dots, f_{i,d}$  can be computed for problem instances  $i \in \mathcal{I}$  coming from this space. Furthermore, let  $\mathcal{A}$  denote a parametrized *target algorithm*, with parameters  $p_1, \dots, p_k$  which may be of categorical or numerical nature. The (finite or infinite) domain of each parameter  $p_i$  is denoted by  $\Theta_i$  such that  $\Theta \subseteq \Theta_1 \times \dots \times \Theta_k$  is the space of all feasible parameter combinations, i.e., the so-called *configuration* or *search space*. A concrete instantiation of the target algorithm  $\mathcal{A}$  with a given configuration  $\theta \in \Theta$  is denoted by  $\mathcal{A}_\theta$ . Furthermore, let  $c : \mathcal{I} \times \Theta \rightarrow \mathbb{R}$  be a cost function from the space of cost functions  $\mathcal{C}$ , which quantifies the cost of running a given problem instance with a given configuration.<sup>2</sup> Depending on the target algorithm,  $c$  may be stochastic and contain noise. Then, ideally, we would like to find the optimal configuration  $\theta^* \in \Theta$  defined as

$$\theta^* \in \arg \min_{\theta \in \Theta} \int_{\mathcal{I}} c(i, \theta) d\mathcal{P}(i) . \tag{1}$$

---

2.  $\mathcal{C}$  merely limits the possible cost functions, but it is needed for formalizing an aggregation function.

However, in practice, the distribution  $\mathcal{P}$  over  $\mathcal{I}$  is unknown, and thus we must resort to solving a proxy problem.<sup>3</sup> To this end, we are provided both a set of training instances  $\mathcal{I}_{train} \subseteq \mathcal{I}$  and an aggregation function  $m : \mathcal{C} \times 2^{\mathcal{I}} \times \Theta \rightarrow \mathbb{R}$ . The aggregation function is usually the arithmetic mean or a variation thereof that is computed over the given problem instances by applying the given configuration to each of them and computing their cost. Similar to empirical risk minimization in machine learning, we then seek to find the configuration minimizing the aggregated costs across the training instances, i.e.,

$$\hat{\theta} \in \arg \min_{\theta \in \Theta} m(c, \mathcal{I}_{train}, \theta). \quad (2)$$

Informally, the problem can be expressed as: given a target algorithm with a set of parameters and a set of problem instances, find a configuration that yields good performance with respect to the cost measure across the set of problem instances. We will refer to automated approaches capable of finding such configurations as (*algorithm*) *configurators*.

**Configuration example** To make AC more accessible and to illustrate the associated challenges, we use the previously mentioned circuit assignment problem with Glucose as a SAT solver. Glucose in version 4 has 41 parameters (with 13 binary and 28 continuous parameter domains) that constitute the configuration space  $\Theta$ . Table 1 shows a subset of the parameters with their values, bounds and the effect they have on the search. Note that we have simplified the parameter descriptions. In fact, the complexity of understanding what the parameters actually do further emphasizes the need for automated AC, as in many cases practitioners may not fully understand the function of the parameters.

Parameter	Type	Domain	Effect
Random Frequency (RF)	Float	[0, 1]	Probability of assigning a variable a random value instead of using a heuristic.
Conflict Factor (CF)	Float	[0, 1]	Factor for comparing the state of the current restart to the state of the search as a whole.
Restart Queue (RQ)	Int	{10, 11, ...}	Moving average window size over clause conflicts for computing a score of the current state.
Preprocessing (PR)	Cat.	{on, off}	Activate or deactivate preprocessing of problem instances before the main search.

Table 1: Subset of Glucose parameters.

Suitable configurations are very problem dependent and need to be chosen carefully for the application at hand. For instance, the *Random Frequency* parameter of Glucose influences the exploration/exploitation tradeoff. A setting of  $p_{RF} = 1$  will lead to a random search. On the contrary,  $p_{RF} = 0$  will let Glucose only assign values based on its heuristic, which could lead to a local optimum. Depending on the search landscape of the circuits, the appropriate value for the random frequency, and all others, may vary. In addition to setting single parameter values, parameter interactions play an important role in algorithm performance. Consider the parameters *Conflict Factor* and *Restart Queue*, which both

3. Formally,  $c$  is also required to be integrable

influence the restart policy (Huang, 2007), in which the search is restarted from the top of the search tree. To force Glucose to perform restarts more often, both  $p_{CF}$  and  $p_{RQ}$  need to be set to reasonable values (Audemard & Simon, 2012). That is, if not set properly, they may contradict each other and result in unintended or ineffective restart behavior.

While a clever practitioner could use his or her domain knowledge to set these parameters manually, usually it is not entirely clear exactly what settings (or parameter interactions) will lead to good performance. Thus, instead of relying on domain knowledge and configuring and testing parameter values manually, a configurator as ParamILS may be used. A practitioner collects problem instances and selects a cost and aggregation function to solve 2. In the context of our example, a specific circuit relates to one problem instance  $i$  for which features  $f_{i,d}$  can be computed as proposed by Kroer and Malitsky (2011). The circuit can be assumed to come from a distribution  $\mathcal{P}$  that specifies all possible circuits for the application and how often specific circuits are expected to occur. The practitioner approximates this distribution by gathering a limited amount of representative circuits from this distribution over time in a training set  $\mathcal{I}_{train}$ . Since the application at hand requires many circuits to be evaluated in a short amount of time, the cost function  $c$  specifies runtime minimization. As an aggregation function  $m$ , the arithmetic mean over the training instances can be used for a given configuration. Based on these inputs, the algorithm configurator ParamILS is then tasked with finding a suitable configuration  $\hat{\theta}$  for Glucose that can quickly solve the example circuits, and should also reduce the runtime for unseen circuits encountered in the future.

## 2.2 Review Scope

We select and review stand-alone AC methods that are suitable to solve the problem described in Section 2.1. To identify relevant contributions in the literature we use the search terms  $\{\textit{Algorithm Configuration}, \textit{Parameter Control}, \textit{Parameter Tuning}\}$  combined with one of  $\{\textit{Automated}, \textit{Automatic}, \textit{Offline}, \textit{Online}, \textit{Realtime}, \textit{Dynamic}, \textit{Instance Specific}, \textit{Per-Instance}, \textit{Multiobjective}, \textit{Feature Based}, \textit{Optimal}\}$  as a prefix. We use Google Scholar as our search engine. We manually filter the search results to only include methods designed to automatically configure solvers without user interaction, and are able to handle large search spaces. Moreover, we omit articles related to algorithm selection (AS) and hyperparameter optimization (HPO). We define these areas in more detail later. We follow citations forwards and backwards from all articles that we accept into the review to find articles that may lack the keywords above. We place such articles into the list of articles and filter them as previously discussed.

**Algorithm Selection (AS)** AS is a sub problem of AC, however, we do not consider it in this work, as it has been considered in several reviews already (Kerschke et al., 2019; Kotthoff, 2016). Thus, we now define what we mean by AS problems so that they can be filtered out of the review. AS can be seen as a special case of instance-specific AC (see Section 8), where the search space contains only one categorical parameter that models the target algorithm choice. In other words, the AS refers to learning how to configure that single categorical parameter, i.e. the algorithm choice, depending on the input instance. Thus it is severely restricted compared to AC. More specifically, the search space in AS is typically small, discrete and consists of a (static) set of algorithms (although new extensions

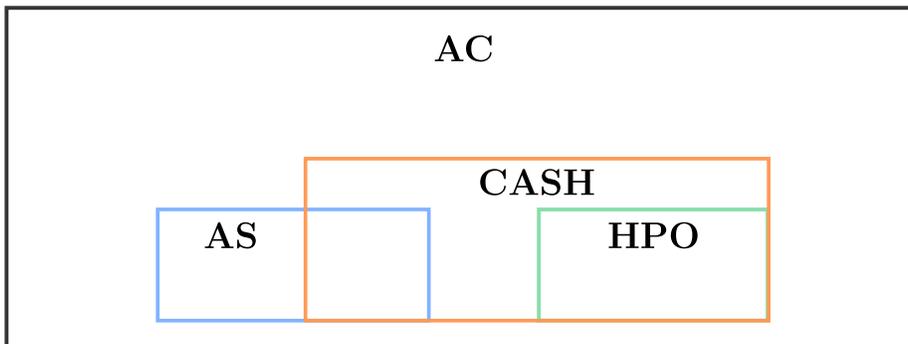


Figure 1: Illustration of the relationship between AC, AS, HPO and CASH.

exist that handle larger spaces (Tornede et al., 2020b)). The search space in AC, in general, is based on the parameters of one target algorithm, and thus, algorithm configurators need to be able to handle much larger, if not even infinite, search spaces (Kerschke et al., 2019).

**Hyperparameter Optimization (HPO)** Our review ignores HPO techniques in addition to AS, as these have also been considered in a number of reviews already (Yu & Zhu, 2020; Luo, 2016; Yang & Shami, 2020; Bischl et al., 2021). Furthermore, before more clearly defining HPO, let us clarify the terminology around the words hyperparameter and parameter. In HPO, parameters that should be set by a user are referred to as hyperparameters, while in AC these are referred to as parameters. HPO refers to hyperparameters since machine learning models usually also contain parameters that are induced from data and are not considered by a configurator. In fact, it is this difference in terminology that leads us to one of the key differences between HPO and general AC, namely that AC methods focus on configuring target algorithms that solve instances of a dataset *independently*, while HPO learns hyperparameters for target algorithms that train parameters on *multiple* instances of a single dataset in tandem.

As HPO is a subset of the AC setting, HPO techniques can, in theory, be used to search for configurations in the general AC setting. In reality, this is seldom done because HPO methods ignore two key functionalities necessary for the general AC setting. First, HPO does not minimize algorithm runtime, which is often performed in general AC. HPO aims at optimizing a solution quality metric, such as predictive accuracy. Of course, AC settings exist where runtime is not a configuration objective, such as when configuring metaheuristics to find the best possible solution in a given time budget. Second, HPO techniques lack a problem instance selection mechanism. Specifically, one configuration in HPO is run on all the instances of one dataset and the result is observed by the configurator. Note that this set should be seen as a single AC problem instance, i.e., the term “instance” is used differently between the HPO and general AC communities. In AC, the configuration needs to be tested on a (sub)set of problem instances before the configurator can infer traits about its quality. Furthermore, HPO can be paired with algorithm selection, which is referred to as the combined algorithm selection and hyperparameter optimization problem (CASH) (Thornton et al., 2013).

**Automated Machine Learning (AutoML, CASH)** The HPO problem can be extended by an algorithm selection component. The task of selecting machine learning algorithms and simultaneously tuning their hyperparameters is formalized by Thornton et al. (2013) as the combined algorithm selection and hyperparameter optimization (CASH) problem. Similar to HPO, the CASH problem can be classified as a sub-problem of algorithm configuration that is restricted to the domain of machine learning. Note that in the setting of AutoML, configurators typically face only a single AC problem instance in the form of a machine learning dataset. Due to this, we do not cover AutoML/CASH within this overview but instead refer the interested reader to comprehensive surveys (Elshawi et al., 2019; Zöllner & Huber, 2021; Hutter et al., 2019).

### 3. Classification

We propose a classification scheme that separately covers (1) the algorithm configuration setting and (2) the configurator itself. More precisely, the *problem view* describes the configuration setting a method tries to address. The *problem view* consists of eight sub-categories with an emphasis on the properties of the problem and the interaction between the configurator and target algorithm. The *configurator view* consists of seven components that portray important aspects of a configurator. Both of these views are interconnected and complementary. Moreover, the configurator view can be interpreted as an answer to a problem setting, where specific features are added to the configurator as a response to the configuration setting. Existing classification schemes proposed in the literature until now (Huang et al., 2019; Eiben & Smit, 2011a,b; Stützle & López-Ibáñez, 2019; Eryoldaş & Durmuşoğlu, 2021) focus solely on the configurator and ignore the problem setting. The proposed taxonomy allows for a description and characterization of methods by aggregating information in tuples. The scheme (especially the *problem view*) can also be used to derive new problem scenarios that have not been addressed before by combining different aspects in previously unseen ways.

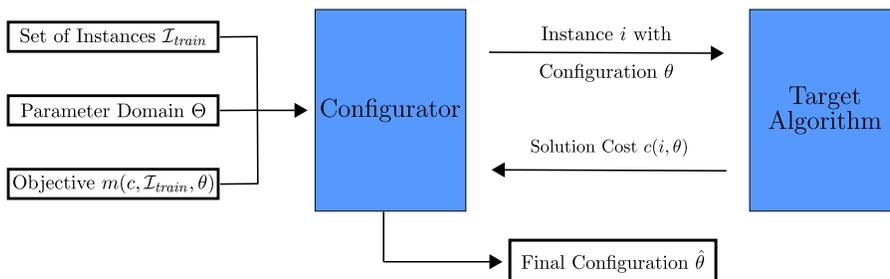


Figure 2: Illustration of offline AC

#### 3.1 Problem View

The components of the *problem view* (Table 2) characterize a problem setting a configurator is meant for and therefore influence the configurator’s design. Figure 2 displays these

interconnections and the communication between target algorithm and configurator, as well as the inputs a configurator receives. Note that, except for the *objective function* and *external runtime setting*, all other aspects are mutually exclusive, meaning that an unambiguous setting for a configurator exists. Furthermore, only the *training setting* and *configuration scope* are independent of the target algorithm. In the following, we elaborate on the dimensions and discuss their implications.

Problem aspects	Options		
Training setting	Offline	Realtime	
Configuration scope	Set	Instance	
Search space	Small discrete	Large discrete	Infinite
Target algorithm objective type	Single-objective	Multi-objective	
Objective function*	Solving time	Accuracy	Memory Usage
Target algorithm observation time	During run	Post termination	
Configuration adjustment	Static	Dynamic	
External runtime setting*	Limited	Infinite	

\* Options not mutually exclusive

Table 2: The problem view classification scheme.

**Training setting** Different training modes for configurators are possible: offline and real-time training or a combination of the two. In the offline setting, the configurator receives the set of training instances  $\mathcal{I}_{train}$  as tuning begins, with which it searches for a suitable configuration. The setting is similar to the classic machine learning training setting, where multiple passes over the training set can be performed by the configurator and sufficient (possibly unlimited) time is available. Model-free and model-based offline methods are outlined in Sections 4 and 5. In the realtime setting, the configurator receives a stream of problem instances, and it should solve each problem instance for the first time with a suitable configuration. While doing so, it can learn from solving the current problem instance to improve the solution time or quality of future configurations. In particular, the configurator is not trained up front, but sequentially during operation, and only one pass over the arriving data is possible. This setting is similar to the online learning setting in machine learning. Realtime methods can be found in Section 7.

**Configuration scope** The configurator can either be required to find a configuration for all problem instances in a set  $\mathcal{I}_{train}$  or for individual problem instances. The former is referred to as a one-fits-all approach, i.e., a single configuration is derived that works well on average over a set of problem instances, whereas the latter requires configurations that are specifically derived for a problem instance. Determining instance specific configurations generally requires the configurator to take into account instance *features*. Instance-specific configurators can be found in Section 8.

**Search space** The problem complexity is highly influenced by the *search space*  $\Theta$  spanned by the target algorithm parameters  $p_1, \dots, p_k$ . In particular, target algorithms may

include a varying number of parameter types, which result in different search space sizes in which parameter transformations (e.g., logarithmic transformation) are possible (Franzin et al., 2018).

**Target algorithm objective type** The configurator may target a single objective or multiple objectives simultaneously.

**Objective function** Different optimization goals exist, so designers must choose an *objective function*  $m(c, \mathcal{I}_{train}, \theta)$  for which a configurator optimizes. Usually, the configurator maximizes a solution quality metric, e.g., predictive accuracy, or minimizes runtime, but other metrics, such as minimizing memory usage, are also conceivable. Note that to handle target algorithm runs where no solution is found, penalty terms (Eggenesperger et al., 2019) can be used.

**Target algorithm observation time** The time a configurator observes the cost  $c(i, \theta)$  returned by the target algorithm can be either at termination of the target algorithm or during its execution.

**Configuration adjustment** Target algorithms provide different opportunities for the configurator to adjust configurations, and this is orthogonal to the *target algorithm observation time*. More precisely, the configuration  $\theta$  can either be adjusted only at the start of a run, where it then remains fixed, or the target algorithm may allow for dynamic adjustments of  $\theta$  during runtime. Configurators that are able to adjust configurations dynamically can be found in Section 10.

**External runtime setting** The target algorithm may be influenced by an externally induced runtime cut off  $\kappa_{max}$ . Many AC settings set  $\kappa_{max}$  explicitly, such as allowing a simulation to only run for a specified amount of time. The configurator can also limit  $\kappa_{max}$ , for example adaptively to reduce runtime, and we refer to this as the *internal runtime setting* as part of the configurator view.

### 3.2 Configurator View

The *configurator view* (Table 3) characterizes algorithm configurators. The scheme does not cover concrete functionalities utilized by configurators such as intensification criteria or creation, selection and elimination of configurations. These functionalities are very difficult to characterize and classify, since for a single mechanism many options with only subtle differences may exist. In the following, we elaborate on relevant directions of the configurator view.

**Solution quality guarantee** Some recent AC methods offer theoretical guarantees regarding the quality of the configuration they return. However, most AC methods are heuristics that offer no proof about configuration quality. Methods that provide guarantees on the solution quality are discussed in Section 6.

**Surrogate models** So-called model-based approaches incorporate surrogate models (Hutter & Hamadi, 2005; Hutter et al., 2006, 2014b) that are able to predict the outcome (e.g., runtime) of a target algorithm run. Such surrogate models (including empirical

Configurator aspect	Setting		
Solution quality guarantee	Heuristic	Proven	
Surrogate models	Model-free	Model-based	
Problem instance features	Featureless	Feature-based	
Target algorithm execution	Sequential	Parallel	
Candidate output	Single configuration	Set configuration	Policy
Configurator objective	Single-objective	Multi-objective	
Internal runtime setting	Limited	Infinite	

Table 3: The configurator view classification scheme.

hardness models) are learned during training and can be used by the configurator to create new configurations, i.e., instead of trying a new configuration directly by running the target algorithm, the surrogate model is used to give a first estimate of the configuration quality. Because the surrogate models use the configuration as input, the models may also be referred to as models over configurations. For the sake of consistency we nevertheless will use the term model-based. Model-free and model-based approaches can be found in Section 4 and Section 5 respectively.

**Problem instance features** For different problem types, features of problem instances can be computed and potentially be used by the configurator. In particular, feature-based approaches use a feature vector  $\mathbf{f}_i$  with problem instance features  $f_{i,1}, \dots, f_{i,j}$ . In the case of the SAT problem, an example for such a feature could be the ratio of the number of variables to the number of clauses (Hutter et al., 2014b). The features are problem family dependent, meaning, e.g., mixed-integer programming problems (MILP) need to be described by different features than SAT problems. Features are commonly used in model-based approaches, e.g., within the surrogate model.

**Target algorithm execution** A configurator may also be characterized by the way it executes and stops configuration runs of the target algorithm. In particular, configurators can run configurations in parallel on multiple cores or sequentially.

**Candidate output** Configurators may return one configuration  $\hat{\theta}$  or a set of configurations  $\hat{\Theta}$ , which can then be used by a decision maker or downstream process. In addition, it is possible for the configurator to return a policy that maps instances to configurations.

**Configurator objective** While most configurators aim at optimizing one single objective  $m(\mathcal{I}, c(i, \theta))$ , some configurators can handle multiple objectives  $M := (m_1, \dots, m_n)$ . Note that the objective here refers to a combination of objective functions related to the configuration itself (e.g., a combination of solution quality and runtime), and not an output vector returned by the target algorithm (e.g., a target algorithm that returns a solution vector consisting of cost and environmental impact). Configurators that are able to handle multiple objective functions are discussed in Section 9.

**Internal runtime setting** In addition to the *external runtime setting*, the configurator may decide to cancel a run of a target algorithm on a given configuration before the

externally set runtime  $\kappa_{max}$  is reached. We refer to the decision made by the configurator to either limit a run by capping or running it until the target algorithm terminates as the *internal runtime setting*. The configurator may cancel runs according to some fixed cut off value or adaptively in relation to the runtime of other configurations. Although AC traditionally focused only on runtime capping, recent work by De Souza et al. (2022) introduces capping mechanisms that can be used with quality metrics. Terminating target algorithm runs before they finish results in censored information for the configurator where only a lower bound for the costs  $c(i, \theta)$  is observed (Hutter et al., 2013; Eggenberger et al., 2018, 2020). Previous work has shown that such right-censored data also needs to be handled properly in similar settings (Tornede et al., 2020c, 2021a; Hanselle et al., 2021).

A full characterization of methods included in this review can be found in Table 4 and 5. Based on the classification scheme, existing configurators and ideas are discussed in the following. Starting with the most basic model-free approaches, methods are grouped by their most important characteristic or the main novelty they introduce.

#### 4. Model-free Methods

We now describe model-free, offline AC approaches that solve the offline AC problem setting described in Section 2.1. Figure 3 shows the general search process that is typically employed by configurators. In this context, there are three design choices that must be addressed by algorithm designers to solve the offline AC problem: (1) the training instance sampling strategy, (2) the creation of configurations (initially and during search), and (3) the evaluation criteria for comparing configurations. Model-free configurators must run the target algorithm with a given configuration to observe its runtime or solution quality. Approaches that estimate the runtime or quality through models are described in Section 5. A variety of methods can be found in the literature that tackle each of the mentioned design choices differently, and we discuss them in the following.

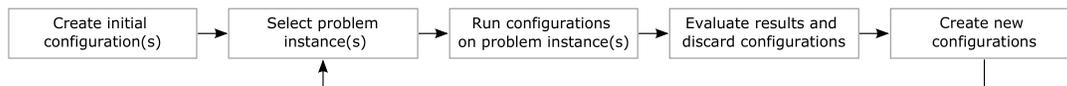


Figure 3: Configuration search process of a configurator

**Calibra** One of the first approaches to systematically search for suitable parameter values is Calibra (Adenso-Diaz & Laguna, 2006). Calibra is a heuristic configurator that finds a configuration for a small discrete space of up to five parameters for a set of problem instances optimizing for solution quality. To do so, Calibra utilizes a combination of factorial experiments and local search.

Factorial experiments are used to compare configurations, and a subsequent local search guides Calibra towards promising configurations. Initial configurations are created by a full factorial experimental design (Fisher, 1937) using the first and third quartile of the value ranges of the parameters as fixed starting points. This factorial design is the limiting

Configurator	Reference	Training setting	Configuration scope	Target algorithm observation time	Configuration adjustment	Target algorithm objective type	Search space
F-Race	Birattari et al. (2002)	O	Set	Pt	Sta	Si	Sd
Calibra	Adenso-Diaz & Laguna (2006)	O	Set	Pt	Sta	Si	Sd
HORA	Barbosa & Senne (2017a)	O	Set	Pt	Sta	Si	Sd
ParamILS	Hutter et al. (2007b)	O	Set	Pt	Sta	Si	Ld
Hydra	Xu et al. (2010)	O	Set	Pt	Sta	Si	Ld
BNT	do Nascimento & Chaves (2020)	O	Set	Pt	Sta	Si	Ld
REVAC	Nannen & Eiben (2006)	O	Set	Pt	Sta	Si	I
Sampling F-Race	Balaprakash et al. (2007)	O	Set	Pt	Sta	Si	I
Iterated F-Race	Balaprakash et al. (2007)	O	Set	Pt	Sta	Si	I
GGA	Ansótegui et al. (2009)	O	Set	Pt	Sta	Si	I
PyDGGA	Ansótegui et al. (2021)	O	Set	Pt	Sta	Si	I
ROAR	Hutter et al. (2011)	O	Set	Pt	Sta	Si	I
SMAC	Hutter et al. (2011)	O	Set	Pt	Sta	Si	I
MBGM	Birattari et al. (2011)	O	Set	Pt	Sta	Si	I
D-SMAC	Hutter et al. (2012)	O	Set	Pt	Sta	Si	I
GGA++	Ansótegui et al. (2015)	O	Set	Pt	Sta	Si	I
irace	López-Ibáñez et al. (2016)	O	Set	Pt	Sta	Si	I
irace(cap)	Cáceres et al. (2017)	O	Set	Pt	Sta	Si	I
SP <sup>1</sup>	Kleinberg et al. (2017)	O	Set	Pt	Sta	Si	I
LeapsAndBounds	Weisz et al. (2018)	O	Set	Pt	Sta	Si	I
WS SMAC <sup>2</sup>	Lindauer & Hutter (2018b)	O	Set	Pt	Sta	Si	I
CapsAndRuns	Weisz et al. (2019)	O	Set	Pt	Sta	Si	I
SP* with Confidence	Kleinberg et al. (2019)	O	Set	Pt	Sta	Si	I
GPS	Pushak & Hoos (2020)	O	Set	Pt	Sta	Si	I
ImpatientCAR <sup>3</sup>	Weisz et al. (2020)	O	Set	Pt	Sta	Si	I
SMAC+PS	Anastacio & Hoos (2020)	O	Set	Pt	Sta	Si	I
S-Race	Zhang et al. (2013)	O	Set	Pt	Sta	Mu	Sd
SPRINTRace	Zhang et al. (2015b)	O	Set	Pt	Sta	Mu	Sd
MO-ParamILS	Blot et al. (2016)	O	Set	Pt	Sta	Mu	Ld
HCRS	Ansótegui et al. (2017)	O	Set	Dr	Dyn	Si	I
DAC-RL	Biedenkapp et al. (2020)	O	Set	Dr	Dyn	Si	I
MATE	Yafrani et al. (2020)	O	In	Pt	Sta	Si	Sd
CluPaTra	Lau & Lo (2011)	O	In	Pt	Sta	Si	Ld
FloTra	Lindawati et al. (2013c)	O	In	Pt	Sta	Si	Ld
SufTra	Lindawati et al. (2013b)	O	In	Pt	Sta	Si	Ld
ISAC	Kadioglu et al. (2010)	O	In	Pt	Sta	Si	I
EISAC	Malitsky et al. (2013a)	O	In	Pt	Sta	Si	I
ISAC++	Ansótegui et al. (2016)	O	In	Pt	Sta	Si	I
PCIT	Liu et al. (2019)	O	In	Pt	Sta	Si	I
ReACT	Fitzgerald et al. (2014)	Rt	In	Pt	Sta	Si	I
ReACTR	Fitzgerald et al. (2015)	Rt	In	Pt	Sta	Si	I
CPPL	El Mesaoudi-Paul et al. (2020b)	Rt	In	Pt	Sta	Si	I

Table 4: Configurators: problem view. <sup>1</sup>Structured Procrastination, <sup>2</sup>Warmstarting SMAC, <sup>3</sup>CapsAndRuns. O: Offline, Rt: Real-time, In: Instance, Pt: Post termination, Dr: During run, Sta: Static, Dyn: Dynamic, Si: Single, Mu: Multi, Sd: Small discrete, Ld: Large discrete, I: Infinite.

Configurator	Reference	Solution guarantee	Surrogate model	Instance features	Algorithm execution	Candidate output	Configurator objective	Internal runtime	Distinguishing feature
F-Race	Birattari et al. (2002)	H	Mf	Fl	S	Si	Si	I	Racing & F-test
Calibra	Adenso-Diaz & Laguna (2006)	H	Mf	Fl	S	Si	Si	I	Taguchi design & local search
Sampling F-Race	Balaprakash et al. (2007)	H	Mf	Fl	S	Si	Si	I	F-Race & sampling
HORA	Barbosa & Senne (2017a)	H	Mf	Fl	S	Si	Si	I	Racing, DOE & local search.
ParamILS	Hutter et al. (2007b)	H	Mf	Fl	S	Si	Si	L	Iterative local search
ROAR	Hutter et al. (2011)	H	Mf	Fl	S	Si	Si	L	Random sampling & racing
S-Race	Zhang et al. (2013)	H	Mf	Fl	S	Set	Mu	I	Sign test & racing
SPINTRace	Zhang et al. (2015b)	H	Mf	Fl	S	Set	Mu	I	S probability test & racing
MO-ParamILS	Blot et al. (2016)	H	Mf	Fl	S	Set	Mu	L	Iterative local search
GGA	Ansótegui et al. (2009)	H	Mf	Fl	Par	Si	Si	L	Genetic algorithm (GA) & racing
PyDGGA	Ansótegui et al. (2021)	H	Mf	Fl	Par	Si	Si	L	Genetic Distributed GGA
ReACT	Fitzgerald et al. (2014)	H	Mf	Fl	Par	Si	Si	L	Pool based & racing
ReACTR	Fitzgerald et al. (2015)	H	Mf	Fl	Par	Si	Si	L	ReACT & Trueskill
REVAC	Nannen & Eiben (2006)	H	Mb	Fl	S	Si	Si	I	Distribution estimation & GA
Iterated F-Race	Balaprakash et al. (2007)	H	Mb	Fl	S	Si	Si	I	F-Race & resampling
MBGM	Birattari et al. (2011)	H	Mb	Fl	S	Si	Si	L	Bayesian Networks
BNT	do Nascimento & Chaves (2020)	H	Mb	Fl	S	Si	Si	L	Bayesian Networks
irace	López-Ibáñez et al. (2016)	H	Mb	Fl	Par	Si	Si	I	F-Race & elitist racing
irace(cap)	Cáceres et al. (2017)	H	Mb	Fl	Par	Si	Si	L	Irace & capping
GPS	Pushak & Hoos (2020)	H	Mb	Fl	Par	Si	Si	L	Golden section search
CluPaTra	Lau & Lo (2011)	H	Mb	Fb	S	Si	Si	I	Clustering trajectories & tuning
FloTra	Lindawati et al. (2013c)	H	Mb	Fb	S	Si	Si	I	CluPaTra & suffix tree encoding
SufTra	Lindawati et al. (2013b)	H	Mb	Fb	S	Si	Si	I	CluPaTra & graph trajectories
SMAC	Hutter et al. (2011)	H	Mb	Fb	S	Si	Si	L	SMBO
HCRS	Ansótegui et al. (2017)	H	Mb	Fb	S	Si	Si	L	GGA & logistic regression
WS SMAC <sup>1</sup>	Lindauer & Hutter (2018b)	H	Mb	Fb	S	Si	Si	L	SMBO & warmstart
DAC-RL	Biedenkapp et al. (2020)	H	Mb	Fb	S	Si	Si	L	Reinforcement learning (RL)
SMAC+PS	Anastacio & Hoos (2020)	H	Mb	Fb	S	Si	Si	L	SMBO & probabilistic sampling
Hydra	Xu et al. (2010)	H	Mb	Fb	S	Set	Si	L	Boosting
D-SMAC	Hutter et al. (2012)	H	Mb	Fb	Par	Si	Si	L	SMBO & configuration queue
ISAC	Kadioglu et al. (2010)	H	Mb	Fb	Par	Si	Si	L	Clustering features & Tuning
EISAC	Malitsky et al. (2013a)	H	Mb	Fb	Par	Si	Si	L	ISAC & retraining
ISAC++	Ansótegui et al. (2016)	H	Mb	Fb	Par	Si	Si	L	ISAC & AS
PCIT	Liu et al. (2019)	H	Mb	Fb	Par	Si	Si	L	Clustering features & Tuning
GGA++	Ansótegui et al. (2015)	H	Mb	Fb	Par	Si	Si	L	GGA & RF
CPPL	El Mesaoudi-Paul et al. (2020b)	H	Mb	Fb	Par	Si	Si	L	ReACTR & bandits
MATE	Yafrani et al. (2020)	H	Mb	Fb	Par	Si	Si	L	Genetic programming & regression
SP <sup>2</sup>	Kleinberg et al. (2017)	P	Mf	Fl	S	Si	Si	L	Emp. mean runtime queuing
LeapsAndBounds	Weisz et al. (2018)	P	Mf	Fl	S	Si	Si	L	Phase-based Bernstein racing (BR)
CapsAndRuns	Weisz et al. (2019)	P	Mf	Fl	S	Si	Si	L	Timeout estimating & BR
SP* with Confidence	Kleinberg et al. (2019)	P	Mf	Fl	S	Si	Si	L	Lower confidence bound SP
ImpatientCAR <sup>3</sup>	Weisz et al. (2020)	P	Mf	Fl	S	Si	Si	L	CapsAndRuns with preprocessing

Table 5: Configurators: configurator view. <sup>1</sup>Warmstarting SMAC, <sup>2</sup>Structured Procrastination, <sup>3</sup>CapsAndRuns. H: Heuristic, P: Proven, Mf: Model-free, Mb: Model-based, Fl: Featureless, Fb: Feature-based, S: Sequential, Par: Parallel, Si: Single, Mu: Multi, I: Infinite, L: Limited.

component of Calibra, as factorial design does not scale. During the local search, Calibra creates new configurations by narrowing down value ranges from the previously tested configurations, successively focusing the search on promising regions of the parameter space. Taguchi  $L_9(3^4)$  design (Roy, 2010) is used to pairwise compare the resulting configurations. The value ranges of the parameters are narrowed in each iteration of a loop until an evaluation budget is exhausted. If a local optimum is found, the result is stored, and the local search is reinitialized.

We note that there are, in fact, some earlier AC approaches that use the design of experiments methodology (DOE). Some of the first approaches for AC based on DOE are given by Coy et al. (2001); Bartz-Beielstein (2003); Ruiz and Maroto (2005) and Ridge and Kudenko (2007), which, like Calibra, in practice do not scale to even moderately sized configuration spaces. However, DOE can be used to preprocess parameters to reduce the search space. Gunawan and Lau (2011) propose using factorial experiments to rank parameters according to their importance and to set unimportant parameters to default values, while using a configurator to optimize the remaining, important ones. In later work, Gunawan et al. (2013) propose to use fractional factorial experiments to decompose parameters into different categories before using the previous procedure to identify important parameters for each of the categories. Finally, Fallahi et al. (2014) use DOE and clustering in an expert system that helps a user derive problem instance specific configurations.

**ParamILS** ParamILS (Hutter et al., 2007b) employs iterated local search (ILS) (Lourenço et al., 2003) to find high quality parameter configurations. It can handle discrete search spaces of large sizes with conditional parameters, and can optimize for either solution quality or runtime, making it one of the first general-purpose algorithm configurators proposed in the literature.

ParamILS pairs a local search with two different types of configuration comparison procedures to determine which configuration to perturb next. More precisely, it uses a one-exchange neighborhood search in which one parameter is changed in each iteration. To avoid local optima, the search is restarted at random. The local search in ParamILS differs from Calibra, where the parameters are changed based on the bounds and midpoint of a narrowing value range. The resulting configurations are compared based on two procedures: BasicILS and FocusedILS. BasicILS compares two configurations on an equal number of problem instance runs, where the same problem instances (and random seeds) are used. A configuration is deemed superior if it provides a better objective value (e.g., mean runtime). In contrast, FocusedILS adjusts the number of evaluations for the configurations in question dynamically, utilizing an approach similar to racing. In particular, it uses a notion of dominance, in which a configuration is accepted over a competing configuration if it has been evaluated on more problem instances and has lower cost on those problem instances.

ParamILS suffers from a few shortcomings that are partly addressed in later work. The main drawback is the fact that it can only handle discrete parameter domains. In addition, it spends a significant amount of time running inferior configurations. To address this, Hutter et al. (2009b) introduce adaptive capping, which terminates configurations as soon as it becomes clear that they will not be better than the current incumbent. Cáceres and Stützle (2017) adjust the one-exchange neighborhood principle of ParamILS to employ

reduced variable neighborhood search (RVNS). This allows for more than one parameter of a configuration to be changed in each local search iteration.

**F-Race** Racing configurators run different parameter configurations against each other and discard inferior configurations based on non-parametric statistical tests. They are inspired by racing mechanisms from the machine learning literature (Maron & Moore, 1993; Moore & Lee, 1994). Birattari et al. (2002) and Birattari (2009) introduce F-Race, the first racing procedure to address AC. F-Race starts by creating configurations to be raced through a full factorial design. Races are then performed sequentially, in which all (remaining) parameter configurations are evaluated on a problem instance. After all runs terminate, F-Race performs a Friedman test to determine whether the results obtained are significantly different from the results of the previous race. If this is the case, F-Race eliminates inferior configurations through pairwise tests. The remaining configurations are raced on the next problem instance until one remains, or a stopping criterion is met.

Extensions to F-Race, such as sampling F-Race, iterated F-Race (I/F-Race) and irace, improve upon the shortcomings of F-Race. Sampling F-Race (Balaprakash et al., 2007; Birattari et al., 2010) creates only a fraction of starting configurations by random sampling, limiting the exponential number of starting configurations. Iterated F-Race (Balaprakash et al., 2007; Birattari et al., 2010) not only shrinks a population of configurations sequentially, but is also able to replenish the set with new configurations between races. To create new configurations, a bivariate normal distribution over the parameter space is used, which in each iteration is parametrized by the previous race winner. irace (López-Ibáñez et al., 2016) is based on iterated F-Race and adds soft-restarts and elitist racing. It replaces the one winner carried over between races by I/F-Race through an elitist set that is kept and updated over races. With the elitist set, configurations must prove viable over a sequence of problem instances, and not win a race just by chance. Soft restarts ensure that configurations in the set do not become too similar over multiple iterations. To this end, a distance-based diversity measure triggers a restart if the population becomes too similar. In this case, diversity is increased by sampling new configurations from a reset probability model. The introduction of adaptive capping (Pérez Cáceres et al., 2017) prevents irace from spending time on evaluating unpromising configurations and makes it competitive for scenarios where costs are related to runtime. Similar to Hutter et al. (2011) and Ansótegui et al. (2015), Cáceres et al. (2017) add a random forest as surrogate to predict the potential costs of new configurations. However, this only leads to minor improvements. van Dijk et al. (2014) propose uRace for deterministic target algorithms.

The principles of racing can also be paired with other configuration approaches. Yuan et al. (2013) combine black box optimizers with a post-selection technique based on racing. Moreover, the optimizer is given a small computational budget for which it finds a set of promising configurations. The configurations are then raced to identify the best configuration among the identified set. Another approach that falls into this line of work proposes to pair evolutionary algorithms (EA) with racing (Yuan & Gallagher, 2005, 2007). The main shortcoming of using gradient-free optimization techniques for the AC problem is that they are limited to continuous parameters, and thus require discrete and categorical parameters to be rounded.

**GGA** The gender-based genetic algorithm (GGA) (Ansótegui et al., 2009) is a population-based approach in which configurations are encoded as individuals in one of two populations: competitive and non-competitive. It supports discrete and continuous parameters and can take parameter interactions supplied by the user into account. It is inherently parallel and optimizes for runtime or solution quality.

GGA combines a strong intensification procedure involving racing individuals from the competitive pool with diversification from the non-competitive population. In each generation of GGA, the competitive population is raced on a random subset of instances that increases linearly with each generation until either a target number of instances is reached or the entire training set is run in each generation. If the number of competitive members of the population is less than the number of CPU cores available on the machine, the race is split into multiple “mini-tournaments” with a number of individuals equal to the number of cores. When configuring for runtime, mini-tournaments are stopped as soon as a set number of individuals in the mini-tournament solve all the instances of the subset; usually one or two individuals. In this way, GGA saves significant runtime without needing to guess an instance timeout that could result in the instance not being solved at all. In the case of quality tuning, the mini-tournaments are run completely and the winners are selected based on the evaluation criteria set by the user.

Once the tournament phase is completed, GGA updates its population using custom recombination and mutation operators. In each generation, one third of the population is replaced and removed from the population. This third is replaced with new individuals that are created by combining a randomly selected winner from the competitive population with a randomly selected non-competitive individual. The recombination operator selects parameters from the competitive and non-competitive individual in a conservative fashion based on a dependency tree that represents how the parameters interact with each other. In subsequent work (Ansótegui et al., 2015), the recombination operator is extended to use a random forest surrogate. A mutation operator is applied to randomly change a small percentage of the parameters using Gaussian distributions for continuous and discrete parameters (with a mean equal to the current value of the parameter) and a uniform distribution for categorical parameters. The new individuals are inserted into the competitive or non-competitive population at random, and GGA moves to the next generation. GGA terminates if it runs out of time or a goal generation is reached and the average performance of the population stops improving. Ansótegui et al. (2021) present PyDGGA, an enhanced, distributed version of GGA, that is optimized for high-performance computing clusters. Further improvements as well as an approach to instance-specific configuration can be found in Ansótegui et al. (2021).

**HORA** The heuristic oriented racing algorithm (HORA) (Barbosa & Senne, 2017a,b) combines elements from racing, DOE and neighborhood search. It can handle small search spaces and is able to optimize for either solution quality or runtime, however no capping is performed during races. HORA utilizes a response surface methodology to find an initial set of good configurations, which are used as seeds to iteratively create and race new configurations. More precisely, a few problem instances are selected at the start from the training set from which initial configurations are found through DOE. The resulting configurations are used as the basis for the subsequent search, in which they are altered using a search

method that is similar to the one exchange neighborhood search of ParamILS. To compare the resulting configurations, races are conducted analogous to F-Race. The Friedman test is used to discard configurations. The cycle of generating new configurations and racing them continues until a stopping criterion is reached.

**ROAR** Hutter et al. (2011) introduce random online aggressive racing (ROAR), an offline configurator that samples new configurations uniformly at random and compares these to the current incumbent using an intensification mechanism that is similar to FocusedILS. In particular, ROAR requires as many target algorithm runs for a new configuration as for the current incumbent. The difference to FocusedILS is that for different, new configurations, the order of the problem instances they are tested on is not fixed. Moreover, problem instances and seeds to test new configurations are sampled at random from the set of evaluations the current incumbent was tested on. While being evaluated on these samples, new configurations are rejected as their cost is revealed to be clearly worse than the current incumbent. A configuration is accepted if its cost is better than the incumbent over all problem instances and seeds considered so far. This principle of aggressive racing is also later used by SMAC, with the addition of a predictive model.

**GPS** Motivated by the belief that parameter landscapes may not be as complex as assumed (see also (Harrison et al., 2019)), Pushak and Hoos (2020) propose Golden Parameter Search (GPS), an offline procedure that exploits simple structures of parameter spaces by searching parameter configurations semi-independently in parallel. In particular, it works based on the assumptions that (1) numeric target algorithm parameters are uni-modal and (2) that there are no strong interactions between most of the parameters.

GPS combines elements of the golden section search algorithm (Kiefer, 1953) with common AC methods such as racing, capping and intensification. The search is based on parameter ranges (so-called brackets) that are believed to contain the best value for the parameter. The parameter values within the brackets are evaluated in parallel and independently of other brackets using a racing procedure with a permutation test to compare runs. After the target algorithm runs, brackets are expanded and shrunk using the golden section search algorithm, where ranges are updated by shifting values. In case of a bracket being updated, the other brackets are informed about the value adjustment of the respective parameter. In addition, GPS uses a bandit approach based on the relative importance of a parameter to prioritize target algorithm runs for more important parameters.

To speed up evaluation, GPS incorporates capping and intensification mechanisms while continuously populating a queue of configurations that should be run in the future. Initial evaluations are performed only on a small set of problem instances, and this set is gradually increased by sampling from the training set when parameter updates are performed. In addition, GPS modifies the bound multiplier of the capping mechanism to be adaptive itself by introducing a dependence on the size of the current problem instance set, leading to a less aggressive capping strategy compared to ParamILS or irace. Lastly, a queue is maintained that dynamically adjusts the amount of instance-configuration combinations that are run to ensure CPU resources are effectively utilized.

## 5. Model-based Offline Methods

In this section, we discuss model-based, offline AC approaches for solving the offline AC problem setting described in Section 2.1. Model-based methods leverage some form of learned model, such as a random forest, that captures knowledge about the performance of configurations as part of the optimization process. In the following, we provide a short background on sequential model-based optimization (SMBO) approaches, then describe the various model-based methods, starting with random forest methods and followed by Bayesian network based methods.

### 5.1 Sequential Model-based Optimization

SMBO approaches model the AC problem as a black box optimization problem. Formally, given a function  $g : \Theta \rightarrow \mathbb{R}$  with input domain  $\Theta \subset \mathbb{R}^d$ , the goal is to find the point optimizing the function  $g$ :

$$\boldsymbol{\theta}^* := \arg \min_{\boldsymbol{\theta} \in \Theta} g(\boldsymbol{\theta}) . \quad (3)$$

The function  $g$  is referred to as a black box (function) since very few assumptions are made about its structure except for the ability to perform evaluations given a point as input. When defining  $\Theta$  as the configuration space and  $g$  as the performance of the algorithm to configure, the mapping from AC to black box optimization appears straight forward, however a key challenge is deciding which instances from  $\mathcal{I}$  to run with a given configuration and how to integrate the incomplete information about the cost into the model.

SMBO approaches iteratively refine a surrogate model  $\hat{g} : \Theta \rightarrow \mathbb{R}$  with the goal of approximating the original  $g$  as best possible based on point-wise evaluations of  $g$ . To this end, they are powered by two main concepts, first and foremost the aforementioned surrogate model used to approximate  $g$  and a so-called acquisition function  $a : \Theta \times \hat{\mathcal{G}} \rightarrow 2^\Theta$ , where  $\hat{\mathcal{G}}$  is the space of all possible surrogate models  $\hat{g}$ . In short, SMBO approaches iteratively leverage the acquisition function  $a$  to generate a set of configurations to evaluate based on the current approximation of  $g$ , i.e., the surrogate model  $\hat{g}$ . The evaluations of the configurations on  $g$  are provided to the surrogate model to improve its approximation quality and, thus, obtain a better idea of the response surface of  $g$ .

One of the assumptions to be imposed on  $g$ , which is particularly important for the problem of AC, is whether  $g$  is a deterministic or stochastic function. Often, in AC we consider randomized algorithms, meaning that  $g$  is stochastic, i.e., the evaluations of  $g$  contain *noise*. Thus, multiple evaluations at the same point can result in different outcomes. Furthermore, the previously mentioned issue of sampling  $\mathcal{I}$  leads to a noisy estimation of each point, as evaluating all instances is usually too expensive until later stages of configuration.

We give a more detailed description of the SMBO framework for AC based on the pseudocode presented in Algorithm 1 in the following. At the start of the optimization, a so-called initial design strategy is used to generate a set of initial configurations to perform an initial training of the surrogate model  $\hat{g}$ . For a further overview of these strategies, we refer to (Santner et al., 2003). Once the surrogate model is fit on these points, the main loop starts and the following steps are repeated until a stopping criterion is met. First, the acquisition function leverages the current surrogate model  $\hat{g}$  to generate a set

of configurations that  $g$  should be evaluated on. The set of configurations returned by the acquisition function is then passed on to the intensification procedure, which decides how many evaluations are performed per configuration provided. Lastly, the surrogate model is updated based on the configurations chosen by the acquisition function and the corresponding evaluation results provided by the intensification strategy.

---

**Algorithm 1** SMBO(configuration space  $\Theta$ , *initial\_design*, *black box function*  $g$ , *surrogate model*  $\hat{g}$ , *acquisition function*  $a$ )

---

```

1:  $\hat{\theta}^* \leftarrow \text{random}(\Theta)$ 
2:  $X_{init} \leftarrow \text{initial\_design}()$ 
3:  $Y_{init} \leftarrow g(X_{init})$ 
4: Fit  $\hat{g}$  to  $(X_{init}, Y_{init})$ 
5: while stopping criterion not met do
6:    $X \leftarrow a(\Theta, \hat{g})$ 
7:    $Y \leftarrow \text{intensification}(X, g)$ 
8:    $\text{update\_incumbent}(\hat{\theta}^*, X, Y)$ 
9:   Update surrogate model  $\hat{g}$  with  $(X, Y)$ 
10: end while
11: return  $\hat{\theta}^*$ 

```

---

**Surrogate models** The task of the surrogate model, sometimes also called response surface or empirical performance model, is to approximate the original though unknown function  $g$ . Several classes of surrogate models have been considered in the literature, most of which are of a probabilistic nature in order to properly capture uncertainty. The most common ones are Gaussian processes (GP) (Rasmussen & Williams, 2006), tree Parzen Estimators (Bergstra et al., 2011) or random forests (Breiman, 2001). Due to their computational complexity, GPs are often rather impractical for AC problems when considering more than a handful of parameters to optimize.

**Acquisition functions** The acquisition function handles the exploration-exploitation dilemma common to optimization techniques. The function must decide to either return configurations from regions of  $g$  that are likely to yield good performance according to  $\hat{g}$  (exploitation), or to return configurations from regions with larger uncertainty (exploration). A common criterion powering the acquisition function is expected improvement (Mockus, 1974; Jones et al., 1998) (EI), which in its simplest form<sup>4</sup> can be defined as

$$\mathbb{E}[I(\boldsymbol{\theta})] = \mathbb{E} \left[ \max \left\{ 0, g(\hat{\boldsymbol{\theta}}^*) - g(\boldsymbol{\theta}) \right\} \right], \quad (4)$$

where  $\hat{\boldsymbol{\theta}}^*$  is the currently best known configuration, i.e., the incumbent. The expectation is required because  $g(\boldsymbol{\theta})$  is unknown at the time of computation and hence, it is a random variable. Accordingly, to compute the EI, a *probabilistic* surrogate model is required. For a detailed overview of different acquisition functions, we refer to (Frazier, 2018).

**Non-general AC SMBO approaches** Several SMBO approaches exist to perform a limited form of AC, i.e., on only a single instance. We include these methods due to their historical importance to the field of AC, as well as because they may inspire new general AC

---

4. Note that for simplicity we assume that  $g$  is deterministic here.

approaches. Most SMBO based AC approaches are based on the idea of sequential kriging meta-modelling (Huang et al., 2006) (SKO) and sequential parameter optimization (SPO) (Bartz-Beielstein et al., 2005), both of which are based on efficient global optimization (Jones et al., 1998). While the latter is a classical approach to black box function optimization using BO, both SPO and SKO constitute extensions to noisy black box functions; an assumption that is much more realistic for AC. However, both of these approaches still have potential drawbacks. Some of these are fixed by SPO<sup>+</sup> (Hutter et al., 2009a), which improves the intensification scheme, and time-bounded SPO (TB-SPO) (Hutter et al., 2010b), which generalizes SPO<sup>+</sup> to work under (potentially tight) time constraints instead of considering the number of function evaluations as a stopping criterion.

## 5.2 General Model-based AC Methods

**SMAC** Sequential model-based optimization for algorithm configuration (SMAC) (Hutter et al., 2011; Lindauer et al., 2021) can be seen as one of the first fully-fledged model-based AC approaches, as it features solutions for many of the limitations of the previously discussed SMBO techniques. SMAC generalizes TB-SPO to perform configuration over multiple problem instances so that it can support categorical parameters and handle tight time constraints.

To support multiple problem instances, SMAC adapts the intensification strategy of TB-SPO to iteratively evaluate configurations on randomly sampled combinations of seeds and problem instances. When doing so, it ensures that configurations are compared only based on a performance estimate computed on the same randomly sampled set of problem instances. Furthermore, SMAC’s surrogate model can generalize across problem instances by incorporating problem instance features. To this end, a surrogate model is learned on the joint problem instance and configuration space to predict the performance of a given configuration on a given problem instance.

As a means to deal with a mixture of categorical and numerical parameters, SMAC employs a random (regression) forest as a surrogate model, which can inherently handle both types of parameters. The probabilistic nature of the surrogate model required for most acquisition functions is preserved by computing the predicted mean performance and the variance of a configuration as the corresponding statistics on the predictions of the individual trees. Moreover, SMAC leverages an adapted acquisition function. In particular, it uses a special form of the EI criterion and solves the EI maximization problem by a multi-start local search procedure that is equipped with special neighborhoods to deal with categorical hyperparameters. In addition to the configurations obtained through the local search procedure, SMAC evaluates the EI on a set of randomly sampled configurations for the purpose of exploration.

Anastacio and Hoos (2020) propose SMAC+PS, which integrates the idea of probabilistic sampling known from irace into SMAC. This enhancement yields improvements over both SMAC and irace in many cases. In particular, Anastacio and Hoos (2020) account for the problem that many of the completely randomly sampled configurations by SMAC often exhibit rather bad performance and thus, their evaluation yields only limited information. To this end, the authors suggest to sample configurations according to a truncated normal distribution centered around the default configuration.

In (Lindauer & Hutter, 2018b) the authors suggest two different strategies to *warmstart* model-based AC approaches and apply their suggestions to SMAC, leading to significant speedups from days to hours of configuration time. The idea underlying warmstarting is to use the evaluations of configurations from previous runs, i.e., on different problem instance sets, to speed up the configuration process in new runs of the configurator on a new set of instances.

Distributed SMAC (Hutter et al., 2012) (D-SMAC) is an extension of SMAC leveraging parallelization to speed up the configuration process. The main idea behind D-SMAC is to parallelize target algorithm runs onto available workers as much as possible. For this purpose, it maintains a queue of target algorithm configuration evaluations to be performed, which is refilled such that the amount  $k$  of available workers can always be used to its maximum. The intensification strategy is thus adapted to only push a set of evaluations to be performed into the queue instead of actually performing the necessary evaluations. While the workers are performing evaluations, a master process keeps track of the queue, selects configurations to be evaluated and updates the surrogate model. Furthermore, the authors replace EI by a parametrized upper confidence bound (UCB) criterion (Jones, 2001), where the parameter controls the degree of exploration. To allow for an efficient generation of solution candidates based on the acquisition function, the master divides the queue into blocks of  $k$  slots, i.e., one for each worker, and generates the corresponding configuration to be evaluated based on  $k$  differently parameterized instantiations of the UCB criterion.

**GGA++** Ansótegui et al. (2015) adapt the model-free AC approach GGA to include a surrogate model. More precisely, the authors use a surrogate model to evaluate the quality of new configurations. They integrate this within a crossover operator and call it genetic engineering. Recall that GGA contains both a competitive and non-competitive population in which winning configurations from the races between members of the competitive population are recombined with individuals from the non-competitive population. To this end, the crossover operator generates individuals according to the parameter tree crossover of the original GGA method and evaluates them using the surrogate. Note that rather than predicting the solution quality or runtime directly, the surrogate predicts the rank the individual would have in a tournament. The individuals with the best ranks are accepted into the population of the next generation in the same way as in GGA.

While the GGA++ surrogate is based on a random forest model, it differs in a key way. The premise of a random forest is to equally approximate the underlying function over the complete input space. In the case of AC, this is undesirable as only the areas of the input space that correspond to high-quality configurations are of interest. Thus, the authors present specialized splitting criteria that focuses on only the best configurations to increase the focus in high-quality regions of the configuration space while sacrificing approximation quality in other regions.

**Graphical models** Birattari et al. propose an approach we refer to as model-based graphical methods (MBGM) (Birattari et al., 2011) leveraging Bayesian networks (BNs) as a probabilistic surrogate model. Each node in such a network has an associated prior probability distribution over the parameter values. Here, both the network structure and the initial probability distributions of the nodes are supposed to be provided in advance. The approach then works as follows. First, an unseen problem instance from the set of

training instances is chosen. Second, new configurations are created by sampling a value for each parameter one after another according to the conditional probability distribution modeled by the respective part of the BN. The order in which parameters are sampled is imposed by any of the topological orderings of the BN such that the sampling adheres to parameter dependencies. Third, the sampled configurations are evaluated on the unseen problem instance and, based on the performance, a subset of them is chosen to be added to a training dataset. Fourth, the probability distributions of the nodes are updated according to the posterior distribution based on the data contained in the training dataset. This process is repeated until a stopping criterion is reached and the best seen configuration is returned.

A key aspect of the approach is the update procedure for the probability distributions of the nodes. This must ensure that the distributions are shifted towards more promising regions to yield good performance. Correspondingly, Birattari et al. (2011) show how this update can be performed efficiently for numerical and categorical parameters and even for networks for parameters of mixed type (categorical and numerical).

**BNT** do Nascimento & Chaves propose a method called Bayesian network tuning (BNT) (do Nascimento & Chaves, 2020), which follows the same idea of leveraging BN for algorithm configuration, but employs a population of configurations. BNT starts by generating and evaluating an initial population of configurations according to a given initial design, and then iterates over the following steps. First, a subset of the current population corresponding to the best performing configurations is selected. Second, on the basis of this subset, a BN is created leveraging special metrics quantifying parameter interdependence. In contrast to (Birattari et al., 2011), BNT does not assume a fixed network structure to be given, but creates it itself in each iteration. Third, new configurations are created by sampling from the BN as described earlier. Finally, the newly created configurations are evaluated and replace the worst configurations in the current population. This process is repeated until a stopping criterion is reached.

**REVAC** Parameter relevance estimation and value calibration (REVAC) (Nannen & Eiben, 2006, 2007) is an offline configurator for continuous parameters. It estimates probability density functions over parameter values, where the distribution can be used to draw conclusions about parameter relevance and parameter ranges after termination. It does not incorporate a capping mechanism, and continuous as well as categorical values need to be discretized.

REVAC is based on the ideas of estimation of distribution algorithms and genetic algorithms. While creating and evaluating configurations, REVAC maintains a probability density distribution for each parameter. More precisely, REVAC starts with a uniform distribution, which is iteratively refined as configurations are created and evaluated, giving higher probabilities to parameter regions that perform well. New configurations are created by sampling from the estimated distribution, where a population approach is used that smoothens the distribution. The state of the probability model after termination is used to estimate the importance of parameters, and the Shannon entropy can be used to estimate the number of evaluations necessary to reach a defined target cost. Smit and Eiben (2009) add racing and a mechanism called sharpening to REVAC that, similar to GGA, gradually increases the number of instances a configuration is tested on. Another approach that

utilizes evolutionary operators is given by Oltean (2005), in which genetic programming is used on problem dependent C-programs to find configurations. The main drawback here is that crossover and mutation are dependent on the problem instance type (e.g., SAT, MILP, etc.).

## 6. Theoretical Guarantees

The field of AC began with heuristics that offered no guarantees as to the quality of the configurations found. Recently, there has been significant progress in providing a theoretical foundation to AC, offering bounds on the quality of the configurations found. First, we consider current results regarding the number of training instances and the structure of the class of cost functions guaranteeing good generalization of algorithm configurators. Then, we discuss and highlight the recent contributions focusing on the theoretical analysis with respect to the worst-case runtime of algorithm configurations.

### 6.1 Generalization Guarantees

Most of the algorithms configurators encountered so far use the empirical mean as the aggregation function  $m$  in (2), which in light of the objective function given by the expected costs in (1) is a natural choice. This choice of aggregation function raises two urgent questions:

1. Given some finite overall computational budget for each algorithm configuration at hand (i.e., a limit on the number of times a configuration can be run in total), how should this budget be distributed across the set of training instances to obtain the most accurate estimate via the empirical mean for (1) for some configuration  $\theta$ ?
2. How many training instances are needed so that an arbitrarily small estimation error of the minimum (1) can be guaranteed with high probability?

**Distributing the computational budget** Birattari (2004) analyzes the first question under the assumption of an infinite set of training instances and a computational budget  $N$  for a configuration  $\theta$ . It is shown that one single run on  $N$  many problem instances leads to the empirical mean estimate with minimal variance. However, the assumption of an infinite set of training instances is obviously quite restrictive for practical applications.

The latter result of Birattari (2004) has been generalized by Liu et al. (2020) to the more relevant scenario, in which the number of training instances is finite. They show that the empirical mean estimate with the minimal variance is obtained by distributing the available computational budget  $N$  of a configuration as evenly as possible across the training instances. More precisely, if  $K$  is the number of training instances, then the number of runs  $n_i$  on a problem instance  $i$  should be such that  $n_i \in \{\lfloor N/K \rfloor, \lceil N/K \rceil\}$ . Note that this distribution scheme is employed by some of the algorithm configurators discussed before, such as ParamILS, irace, and SMAC.

**Probably approximately correct (PAC) learning** The estimation error of the empirical mean estimate is usually set to be its absolute deviation from its population counterpart,

i.e., for any  $\boldsymbol{\theta} \in \Theta$  it is given by

$$\left| \frac{1}{|\mathcal{I}_{train}|} \sum_{i \in \mathcal{I}_{train}} c(i, \boldsymbol{\theta}) - \int_{\mathcal{I}} c(i, \boldsymbol{\theta}) d\mathcal{P}(i) \right|. \quad (5)$$

With this in mind, the general approach to answer the second question is to derive high probability bounds on the estimation error (5) holding for any configuration  $\boldsymbol{\theta} \in \Theta$ , which are monotonically decreasing with the number of training instances, i.e.,  $|\mathcal{I}_{train}|$ . Given a desired value of the estimation error, say  $\varepsilon > 0$ , one can obtain an answer to the second question by equating the derived bound and the estimation error, and then solve this equation with respect to the number of training instances.

Needless to say, these bounds will highly depend on the desired high probability level, the class of cost functions  $\mathcal{C}_\Theta = \{i \mapsto c(i, \boldsymbol{\theta}) \mid \boldsymbol{\theta} \in \Theta\}$  or its dual function class  $\mathcal{C}_\Theta^* = \{\boldsymbol{\theta} \mapsto c(i, \boldsymbol{\theta}) \mid i \in \mathcal{I}\}$ , as well as on the distribution of the problem instances  $\mathcal{P}$  or the set of training instances  $\mathcal{I}_{train}$ . If the latter dependency is taken into account the high probability bounds or equivalently the resulting guarantees for the number of training instances are said to be *data-dependent*, and otherwise *worst case* (aka *distribution-free*) high probability bounds/guarantees.

Liu et al. (2020) derive worst case high probability bounds on the estimation error (5) for any  $\boldsymbol{\theta} \in \Theta$ , if the configuration space  $\Theta$  is finite and the cost function takes values only in some compact interval. Assuming that the functions in the dual function class  $\mathcal{C}_\Theta^*$  are all Lipschitz continuous, the authors show high probability bounds on the estimation error if the configuration space is infinite, but bounded.

Guided by the observation that the shape of the cost functions in  $\mathcal{C}_\Theta$  has a piecewise structure in many domains, i.e., piecewise-constant/linear, etc., Balcan et al. (2019) derive worst case high probability bounds on the estimation error (5) for any  $\boldsymbol{\theta} \in \Theta$  for classes of cost functions exhibiting this structure. Such piecewise structures of the cost function have been observed for branch-and-bound AC problems (Balcan et al., 2017) or linkage-based hierarchical clustering algorithms (Balcan et al., 2018, 2020a). In a subsequent work, Balcan et al. (2020c) generalize their results by replacing the piecewise-structural assumption on  $\mathcal{C}_\Theta$  by an  $L^\infty$ -norm approximation assumption of  $\mathcal{C}_\Theta^*$ . More precisely, it is assumed that any function in  $\mathcal{C}_\Theta^*$  can be approximated uniformly over the configuration space  $\Theta$  by a function from a class of functions having small Rademacher complexity<sup>5</sup>. Unlike the previous worst case high probability bounds, the authors derive data-dependent high probability bounds based on the empirical Rademacher complexity of the class of functions approximating  $\mathcal{C}_\Theta$  with respect to the  $L^\infty$ -norm. Roughly speaking, the result takes advantage of the fact that the Rademacher complexity of  $\mathcal{C}$  is small if the Rademacher complexity of the class of functions approximating its dual function class  $\mathcal{C}_\Theta$  is small. However, it is shown that the latter fact relies heavily on the approximation with respect to the  $L^\infty$ -norm, as it is also shown that it is impossible to derive non-trivial generalization guarantees if the approximation holds only under the  $L^p$ -norm with  $p < \infty$ . This is due to the fact that a small Rademacher complexity of the class of functions approximating  $\mathcal{C}_\Theta$  with respect to the  $L^p$ -norm with  $p < \infty$  does not imply a small Rademacher complexity of  $\mathcal{C}$ .

---

5. Rademacher complexity (Bartlett & Mendelson, 2003) is a commonly used quantity in statistics and machine learning to measure the complexity of a function class.

## 6.2 Runtime Analysis

Initiated by the work of Kleinberg et al. (2017), the AC problem has recently attracted a great deal of research focusing on theoretically grounded approaches regarding the design and motivation of a configurator, if the relevant cost function  $c$  (see Section 2.1) considered is the runtime of configuration  $\theta$  on problem instance  $i$ . These approaches take into account, on the one hand, how close the supposed best algorithm configuration returned by the configurator is to the actual optimal configuration and, on the other hand, how long it takes on average to return it in the worst case. For this, a definition of an optimal algorithm configuration is required, as is a measure of closeness. Assuming that problem instances are drawn with some distribution  $\mathcal{P}$  from  $\mathcal{I}$ , the expected runtime of a configuration  $\theta$  is  $R(\theta) = \mathbb{E}_{i \sim \mathcal{P}}(c(i, \theta))$ . Thus, the configuration with the optimal expected runtime is given by  $\operatorname{argmin}_{\theta \in \Theta} R(\theta)$  having (optimal) expected runtime

$$\text{OPT} := \inf_{\theta \in \Theta} R(\theta).$$

The search for the optimal configuration is generally too ambitious, as the total runtime required for the configurator must be extraordinarily large (possibly infinite) to guarantee that the best algorithm configuration returned by the configurator is in fact the optimal one with high probability.

As a workaround, one can leverage the idea underlying PAC learning (Valiant, 1984) to the problem at hand. The basic idea is to relax the goal of finding the optimal configuration itself and, instead, find a configuration that is considered to be “good enough”. As there are potentially several such “good enough” configurations<sup>6</sup>, this relaxation of the goal allows the search to be completed in less (and, thus, feasible) time. In this context, “good enough” means that the expected runtime is only worse than the optimal expected runtime up to a multiplicative factor of  $1 + \varepsilon$  for some fixed precision parameter  $\varepsilon > 0$ . Formally, a configuration is said to be  $\varepsilon$ -optimal (“good enough”) iff

$$\mathbb{E}_{i \sim \mathcal{P}}(c(i, \theta)) \leq (1 + \varepsilon)\text{OPT}.$$

However, this relaxation of the target is problematic in the context of AC problems, since the runtimes of configurations often exhibit a heavy-tailed distribution. Indeed, it is not difficult to construct an example based on such distributions in which any (sensible) configurator would, in the worst case, take infinitely long to find an  $\varepsilon$ -optimal configuration; see for instance (Vitercik, 2021, Example 11.1.1).

In light of the prevalence of heavy-tailed distributions in the realm of AC problems, the remedy is to run configurations only up to some timeout  $\kappa \geq 0$  chosen by the configurator that could vary over the runs. In practice, this means that one observes  $\min(c(i, \theta), \kappa)$  if configuration  $\theta$  is run on  $i$  with timeout  $\kappa$ , and also whether the configuration runs into a timeout or solves the problem instance within  $\kappa$ . This gives rise to the  $\kappa$ -capped expected runtimes of a configuration  $\theta$  defined by  $R_\kappa(\theta) = \mathbb{E}_{i \sim \mathcal{P}}(\min(c(i, \theta), \kappa))$ , which now take the place of the uncapped runtimes  $R(\theta)$ <sup>7</sup>. However, the introduction of timeouts inevitably means that a certain proportion of problem instances may not be solved by a

6. Of course, the optimal configuration itself is always “good enough”.

7. Technically,  $R(\theta) = R_\infty(\theta)$ .

configuration within the given timeout. This immediately raises the question of how to choose the timeouts so that the proportion of unresolved problem instances is tolerable, but at the same time, the ( $\kappa$ -capped) expected runtime is not too far from the optimal one? The notion of  $(\varepsilon, \delta)$ -optimality of a configuration is introduced by Kleinberg et al. (2017) and provides an intuitive answer. A configuration is  $(\varepsilon, \delta)$ -optimal if there is a timeout  $\kappa \geq 0$  at least as large as the  $\delta$ -quantile of the configuration's runtime distribution, such that the configuration's  $\kappa$ -capped expected runtime is at most  $(1 + \varepsilon)\text{OPT}$ . Formally,

$$\boldsymbol{\theta} \text{ is } (\varepsilon, \delta)\text{-optimal} \iff \exists \kappa \geq 0 : R_\kappa(\boldsymbol{\theta}) \leq (1 + \varepsilon)\text{OPT} \wedge \mathbb{P}_{i \sim \mathcal{P}}(c(i, \boldsymbol{\theta}) > \kappa) \leq \delta. \quad (6)$$

This notion of  $(\varepsilon, \delta)$ -optimality has been adopted by subsequent works and slightly modified in some cases. Nonetheless, the core idea remains the same, namely that one is willing to leave a fixed share of problem instances unsolved (captured via the  $\delta$ -quantile) to improve the expected runtime, such that it is close to optimal (captured by the additive term  $\varepsilon\text{OPT}$ ). Kleinberg et al. (2017) show that the worst-case expected runtime of any configurator to return an  $(\varepsilon, \delta)$ -optimal configuration (with probability of at least  $1/2$ ) is of order  $\Omega(\frac{|\Theta|}{\delta\varepsilon^2}\text{OPT})$ , when the number of configurations is finite.

**From finite to infinite number of configurations.** To obtain a lower bound result for the case of a large or even possibly (uncountable) infinite number of configurations, the notion of  $(\varepsilon, \delta)$ -optimality needs to be modified, as it is the main limiting factor in this regard. Indeed, due to its definition in (6), any optimal configurator is forced to explicitly consider each configuration to decide on its  $(\varepsilon, \delta)$ -(sub-)optimality. To this end, in a scenario with a large number of configurations, the goal is relaxed to find a  $(\varepsilon, \delta)$ -optimal configuration after excluding the  $\gamma \in (0, 1)$  fraction of best configurations from  $\Theta$  with respect to the expected runtime. Formally, let  $\text{OPT}_\gamma = \inf_{x \in \mathbb{R}_+} \{x \mid \mathbb{P}_{\boldsymbol{\theta} \sim \text{Unif}(\Theta)}(R(\boldsymbol{\theta}) \leq x) \geq \gamma\}$  be the optimal expected runtime after excluding the  $\gamma \in (0, 1)$  fraction of best configurations, then a configuration  $\boldsymbol{\theta}$  is  $(\varepsilon, \delta, \gamma)$ -optimal iff

$$\exists \kappa \geq 0 : R_\kappa(\boldsymbol{\theta}) \leq (1 + \varepsilon)\text{OPT}_\gamma \wedge \mathbb{P}_{i \sim \mathcal{P}}(c(i, \boldsymbol{\theta}) > \kappa) \leq \delta. \quad (7)$$

Kleinberg et al. (2017) provide a result on the worst-case expected runtime of any configurator to return an  $(\varepsilon, \delta, \gamma)$ -optimal configuration for specific choices of  $\delta$  and  $\gamma$ , which is quite similar to the finite case by replacing  $\text{OPT}$  by  $\text{OPT}_\gamma$  and  $|\Theta|/\delta$  by a parameter common to the choices of  $\delta$  and  $\gamma$ .

Another convenient tool provided by the authors is how to turn a configurator with theoretical guarantees for finding  $(\varepsilon, \delta)$ -optimal configurations for a finite configuration space to finding  $(\varepsilon, \delta, \gamma)$ -optimal configurations for an infinite configuration space. This can be done via uniform sampling from  $\Theta$  as follows. If a configuration is sampled uniformly at random from  $\Theta$ , then obviously this configuration belongs with probability  $\gamma$  to the  $\gamma$  proportion of the best configurations, i.e., the best  $\gamma$  configurations. Thus, if  $n$  many configurations are sampled uniformly at random from  $\Theta$ , then the probability that at least one among these  $n$  many belongs to the best  $\gamma$  configurations is at least  $1 - (1 - \gamma)^n$ . Now, fixing a failure probability  $\zeta \in (0, 1)$  for the non-occurrence of the latter event, one can solve the resulting inequality with respect to  $n$  to obtain that  $\lceil \frac{\log(\zeta)}{\log(1-\gamma)} \rceil$  samples are sufficient to guarantee that at least one configuration within the sample belongs to the best  $\gamma$  configurations with probability at least  $1 - \zeta$ .

Balcan et al. (2020b) provide an alternative way to the uniform sampling approach to obtain a finite set, which includes at least one configuration with “good enough” performance. Here, “good enough” performance is again in terms of  $(\varepsilon, \delta)$ -optimality of a configuration, which, however, is defined in a slightly different way as in (6), namely

$$\boldsymbol{\theta} \text{ is } (\varepsilon, \delta)\text{-optimal} \iff R_{t_{\boldsymbol{\theta}}(\delta)}(\boldsymbol{\theta}) \leq (1 + \varepsilon)\text{OPT}_{\alpha\delta}, \quad (8)$$

where  $\text{OPT}_{\alpha\delta} = \min_{\boldsymbol{\theta}} R_{t_{\boldsymbol{\theta}}(\alpha\delta)}(\boldsymbol{\theta})$ ,  $t_{\boldsymbol{\theta}}(x)$  is the  $x$ -quantile of  $\boldsymbol{\theta}$ ’s runtime distribution and  $\alpha \in (0, 1)$  is a slack parameter that is introduced for details outside the scope of this work<sup>8</sup>. It is worth noting that the authors allow other cost functions than runtime in their work, as the notion of  $(\varepsilon, \delta)$ -optimality can also be used for such variants. Moreover, the theoretical guarantees as well as the design of their approach assume that the cost functions  $\mathcal{C}_{\Theta}$  (cf. Subsection 6.1) are piecewise constant.

**Structured procrastination (SP)** Kleinberg et al. (2017) propose an AC technique in which the idea is to postpone potentially hard problem instances and solve supposedly easier problem instances first. It thus only spends time on hard instances if it is unable to solve easy ones, reducing the overall time the technique needs. The authors prove that SP returns an  $(\varepsilon, \delta)$ -optimal configuration with high probability, and that the runtime is optimal up to a logarithmic factor.

We briefly describe how SP works for large parameter spaces, as this is the more realistic case of the approach for practical applications. The approach uses a double-ended queue  $Q_{\theta}$  that stores (instance, timeout) pairs for each configuration. A pair is pulled from the front of the queue and run. Should the configuration not finish the instance within the timeout, the instance is placed at the back of the configuration’s queue with double the timeout. If the instance is completed within the timeout, this information is saved and the instance is not considered again for configuration  $\theta$ . In each iteration of the approach, the  $Q_{\theta}$  is chosen with a  $\theta$  with minimal average performance observed so far. To avoid requiring the user to specify  $\delta$ , SP is implemented as an anytime algorithm that reduces  $\delta$  over the course of its execution. Note that SP returns the configuration with the longest total execution time rather than the best empirical mean due to theoretical reasons.

**LeapsAndBounds** Weisz et al. (2018) propose a phase-based algorithm configurator called LeapsAndBounds that tries to guess an appropriate upper bound on the optimal runtime in each phase by doubling the guess of the bound after each failed phase. The authors provide an upper bound on the worst-case total amount of time their method needs for finding an  $(\varepsilon, \delta)$ -optimal configuration with high probability, which improves upon the result of SP. Furthermore, the superiority of LeapsAndBounds over SP is confirmed empirically on a benchmark of SAT solvers.

In each phase of LeapsAndBounds, each configuration is given a time budget, and a number of problem instances to be solved within the budget. Instances are chosen depending on the phases that have already passed, with specific choices of the phase-dependent quantities being based on empirical Bernstein stopping (Mnih et al., 2008). This mechanism takes the range and the empirical variance of capped runtime observations into account. If a configuration exhausts the time budget without having solved all problem instances, its

---

8. This alternative notion of  $(\varepsilon, \delta)$ -optimality is due to Weisz et al. (2019).

expected runtime is above the phase-dependent upper bound guess with high probability. If this happens for all configurations, the phase has failed and the next phase is started. In the case that some configurations do not manage to use the time budget completely, the empirical mean of the runtimes is accepted as a suitable estimator for the expected runtime, and the configuration with the lowest mean is returned as an  $(\varepsilon, \delta)$ -optimal configuration. Note that even though the time budget is not exhausted, there may still be instances that hit their timeouts.

**Structured procrastination with confidence** Kleinberg et al. (2019) revises the SP configurator and modifies the selection and return criterion used to choose the configuration to recommend at each step to improve its practical performance while maintaining satisfactory theoretical guarantees. More precisely, the authors derive lower confidence bounds on the expected mean runtime of a configuration, which then take over the role of the empirical mean runtimes in the selection process<sup>9</sup>. Thanks to a careful derivation of these lower confidence bounds, the choice mechanism adapts to the difficulty of the configuration problem such that poorly performing configurations are excluded quite quickly; a crucial property that SP, in general, lacks. In addition, the lower confidence bounds, as well as the minimal length of a configuration’s queue, are chosen such that  $\epsilon$  and  $\delta$  do not need to be specified a priori, which, however, is required by all methods discussed before. Another change to SP is that the configuration returned is the one having the highest number of completed as well as pending tasks in its queue. The rationale behind this modification is that more promising configurations run faster on average and thus have a larger number of completed and pending tasks. The resulting modified version of SP is called SP with confidence (SPC) due to the usage of the (lower) confidence bounds.

The authors show the correctness of SPC and derive an improved runtime bound over SP. In addition, they show in an experimental study for a simple benchmark set of SAT solvers that SPC finds reasonable configurations after a smaller amount of computation time than SP and LeapsAndBounds.

**CapsAndRuns** As a follow-up to LeapsAndBounds, (Weisz et al., 2019) propose CapsAndRuns, which improves the former both theoretically and practically. It is also based on Bernstein stopping (Mnih et al., 2008), but extends LeapsAndBounds in how it estimates timeouts. Furthermore, CapsAndRuns refines the upper bound on the total time needed for configuration and adjusts the notion of  $(\varepsilon, \delta)$ -optimality as specified in (8). More precisely, CapsAndRuns first estimates a timeout for each configuration and performs a Bernstein race over the configurations afterward.

CapsAndRuns works in two phases, but note that these phases are different from the phases in LeapsAndBounds. First, the method estimates a timeout for each configuration such that only a  $\delta$ -quantile of the instances exceed the timeout, followed by a phase using this estimate to return an expected runtime estimate using the estimated timeout. To make the search process even more efficient, a global estimator for the optimal expected runtime (considering timeouts) is used for all configurations across both phases to eliminate suboptimal configurations as early as possible. This elimination happens either directly after the first phase, when it turns out that the configuration is too slow, or during the

---

9. This modified selection rule adopts the *optimism in the face of uncertainty principle*, which is a popular paradigm in the realm of reinforcement and bandit learning problems (Lattimore & Szepesvári, 2020)

second phase, as soon as one is confident enough that the configuration’s empirical mean runtime is above the optimal one. A further aspect of the approach is that it can be parallelized across configurations.

In other words, the (sub-)optimality of configurations is measured by means of the closeness of their expected runtimes capped at their  $\delta$ -quantile to the optimal expected  $\alpha\delta$ -quantile-capped runtime.

**ImpatientCapsAndRuns** Guided by the observation that heuristic configuration approaches achieve appealing practical performance by quickly discarding less promising configurations based on only a few observations of their runtimes, Weisz et al. (2020) propose the ImpatientCapsAndRuns (ICAR) algorithm, which builds on CapAndRuns with a more aggressive elimination strategy. This is achieved through a preprocessing mechanism for filtering configurations that are unlikely to be optimal that is run before entering CapsAndRuns. Weisz et al. (2020) prove that ICAR finds an  $(\varepsilon, \delta, \gamma)$ -optimal configuration with high probability for specific ranges of  $\varepsilon, \delta$  and  $\gamma$ .

The preprocessing technique in ICAR is essentially a stripped down version of the two phases of CAR, except that the internal statistics for elimination are chosen differently. Note that this still depends on the global estimate of the optimal expected runtime (with timeouts). For this reason, the preprocessing and the subsequent two phases of CAR are executed successively one after the other on an ever-increasing pool of configurations (batches), with configurations that have already been eliminated no longer being taken into account. This ensures that the global estimator gradually improves, which, in turn, leads to the pre-elimination by the preprocessing routine becoming more and more precise and aggressive.

As a byproduct of the theoretical analysis of ICAR, the authors improve the CAR algorithm by refining the choice of one of its key internal statistics. Furthermore, experimental studies show that ICAR has significantly better practical performance for three benchmark datasets compared to the original CAR and its improved version.

**ParamRLS with RLS<sub>k</sub>** Since the ParamILS algorithm (see Section 4) does not provide any theoretical grounding, Hall et al. (2019) introduce ParamRLS, which is a simplified version of ParamILS. The main difference is that ParamRLS uses a random local search excluding restarts instead of an iterated local search. This modification allows the authors to analyze theoretically the impact of the timeout on the expected number of required configuration evaluations. They prove that at least a timeout of  $\Omega(n \log n)$  is necessary for a problem of size  $n$  while tuning the target algorithm RLS<sub>k</sub> considering the configuration time for the OneMax problem function class. Moreover, they show that at least a timeout of  $\Omega(n^2)$  is needed for the Ridge\* problem function class. In addition, they identify  $k = 1$  as optimal for the OneMax function class when optimizing the fitness values of the configurations.

Specifically, ParamRLS initializes the configurator randomly and then increases or decreases a single parameter chosen uniformly at random by 1 or by 2, also uniformly at random, in each time step. This step is repeated until no parameter change yields an improvement anymore. Two different variants are considered. First, ParamRLS-T, in which the target metric is the optimization time, and second, ParamRLS-F, which identifies the configuration that yields the best-found fitness value. The local search algorithm RLS<sub>k</sub> has

only one parameter  $k$ , which gives the number of bits in the configuration that are flipped in each iteration of the search.

**ParamRLS with (1+1)EA** After the successful theoretical analysis of the simple random local search, Hall et al. (2020a) consider a more complex AC scenario by tuning the mutation rate  $\chi$  of the target algorithm  $(1+1)_\chi$ EA with ParamRLS.  $(1+1)_\chi$ EA works by flipping each of the  $n$  bits of the current configuration independently with probability  $\chi/n$ . Once again, the authors analyze the required timeouts for the runtime metric and the best found fitness value on the two benchmark problem classes Ridge and LeadingOnes. They further prove that all configurators which are using the runtime as a performance metric require a timeout at least as large as the expected time to identify the optimal configuration. Thus, this problem scenario needs larger timeouts than in the case of the best fitness performance metric.

**Harmonic mutation operator** Inspired by the insights from the above analyses of ParamRLS, Hall et al. (2020b) design a harmonic mutation operator for the configurations that provably leads to faster performance in the case of single parameter target algorithms. In fact, the authors prove that it tunes single-parameter algorithms in polylogarithmic time for (approximately) unimodal parameter spaces. Even in the worst case, the harmonic mutation operator only slows down the algorithm by at most a logarithmic factor. In an experimental analysis, the harmonic mutation operator is shown to be superior to the  $l$ -step and random mutation operators used for ParamRLS and ParamILS.

To be more precise, the harmonic mutation operator selects a parameter uniformly at random and samples a step size according to the harmonic distribution. In fact, the probability for a step size  $d$  is  $1/(d \times H_{\Phi-1})$  with  $H_m$  as the  $m$ -th harmonic number  $H_m = \sum_{k=1}^m \frac{1}{k}$  and  $\Phi$  as the range of possible parameter values. Then the best parameter value at distance  $\pm d$  is returned.

## 7. Realtime Methods

Realtime methods relax the assumption made by offline methods that a representative training set is available since, in reality, such training sets might not always be available or costly to obtain. Due to possibly changing business models and requirements, they further allow  $\mathcal{P}_t$ , and therefore problem instances that are drawn from it, to change over time. This phenomenon is equivalent to concept drift in machine learning (Gama et al., 2014). Ultimately, this may lead to a growing disparity between  $\hat{\theta}$  and  $\theta^*$ , possibly resulting in diminishing performance of the offline tuned configuration in production. Realtime algorithm configurators have been developed that can provide configurations on an ongoing, per-instance basis. That is, the training set  $\mathcal{I}_{train}$  is not needed upfront, but problem instances are solved sequentially as they arrive and are used for configuration adjustments on the fly. In terms of their design, realtime configurators incorporate the same components discussed for offline configurators.

For the sake of completeness, we formulate the realtime configuration problem, which sometimes is also referred to as online tuning or parameter control, more formally. We are interested in a configurator  $o : \mathcal{H} \times \mathcal{I} \rightarrow \Theta$  that for a given time step  $t$  and the corresponding problem instance  $i_t \sim \mathcal{P}_t$  provides a configuration  $\theta$  based on the history  $h_t$ . The history  $h_t$

consists of triplets of the form  $\{(i_k, \boldsymbol{\theta}_k, c_k)\}_{k=1}^{t-1}$  containing the problem instances encountered so far, the configurations used and the resulting cost. The goal is to minimize the (average) cost, which for a final time horizon is given by:  $\mathcal{M} = \frac{1}{T} \sum_{t=1}^T c(i_t, \boldsymbol{\theta}_t)$ . For a new problem instance, the optimal configuration in light of the instances encountered so far is then defined as

$$\hat{\boldsymbol{\theta}}^* \in \arg \min_{\boldsymbol{\theta} \in \Theta} \mathbb{E}[c(i_t, \boldsymbol{\theta}) | h_t] . \quad (9)$$

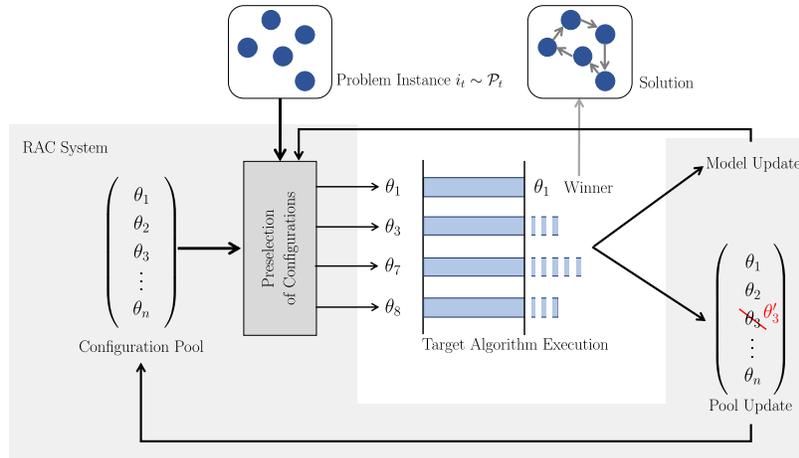


Figure 4: In realtime AC instances arrive and are solved sequentially

**ReACT and ReACTR** The first work to tackle the realtime configuration problem is the method realtime algorithm configuration through tournaments (ReACT) (Fitzgerald et al., 2014). ReACT is a model-free configurator that leverages tournaments of configurations run in parallel for each incoming instance. It maintains a population of configurations and adjusts this in each iteration. It can handle continuous, discrete, and categorical variables. The extension ReACTR (Fitzgerald et al., 2015) introduces a ranking mechanism based on TrueSkill (Herbrich et al., 2006) to rank configurations. This allows for a population that is larger than the number of CPU cores available, leading to more diversity.

While ReACT and ReACTR share tournaments, they differ in population size, configuration elimination and generation mechanisms. Similar to GGA (Ansótegui et al., 2009) both ReACT and ReACTR use parallel tournaments to evaluate individuals and assume that the configuration objective is runtime minimization. Thus, as soon as one of the target algorithm runs finishes, the finisher is declared the winner and all other runs are terminated. While ReACT only allows for a population size equal to the number of CPU cores, ReACTR allows for a bigger population. To avoid having to run multiple tournaments like GGA, ReACTR only runs the top-ranked configurations on the available cores, limiting the exploration of different configurations. Based on the information obtained through the tournament, weak configurations are removed from the pool. ReACT removes configurations based on domination, where it is ensured that a configuration was given enough opportunity to prove its worthiness by requiring another configuration to beat it at least  $m$  times. Individuals are then replenished by random sampling. ReACTR, in contrast, uses its

TrueSkill ranking mechanism to determine which individuals to remove. In addition, it supplements the random sampling of ReACT for making new configurations with a crossover and mutation procedure between the top-ranked individuals. Thus, it can more effectively intensify its search for good configurations around those that have worked well in the past, similarly to I/F-Race or GGA.

**CPPL** The contextual preselection with Plackett-Luce (CPPL) algorithm introduced by El Mesaoudi-Paul et al. (2020b) uses preselection bandits (Bengs & Hüllermeier, 2020) to rank, choose and generate configurations while building on the racing and parallel execution principles from ReACTR (Fitzgerald et al., 2015). It is a model-based, instance-specific configurator that works under the assumption of a Plackett-Luce model (Plackett, 1975; Luce, 2012). Configurations are interpreted as slot machines (“bandits”) that have arms that can be “pulled” by running a configuration to observe their quality. Unlike the classical multi-armed bandit setting (Lattimore & Szepesvári, 2020), where one arm is pulled resulting in a numerical observation (i.e., what is the quality of the pulled arm?), in the preselection bandit setting it is allowed to pull more than one arm at a time leading to qualitative feedback, such as winner feedback (i.e., which of the pulled arms had the highest quality?) or (partial) ranking feedback. Thus, it is an extension of the dueling bandit setting (Bengs et al., 2021) where only two arms are pulled (or dueled against each other). The contextual preselection bandit extension (El Mesaoudi-Paul et al., 2020a) allows CPPL to take problem instance features as additional information into account. That is, the features of a problem instance provide CPPL with information on which arm (configuration) possibly performs best and therefore should be “pulled”.

CPPL builds on ReACTR, but uses the bandit model for key operations. More precisely, the bandit model selects a set of configurations that are to be raced against each other, replacing the TrueSkill ranking mechanism of ReACTR. Based on the obtained (censored) winner feedback, the model is updated using stochastic gradient descent. While ReACTR prunes configurations from the pool using TrueSkill, CPPL uses the upper confidence bounds on the estimated performance of the configurations for pruning. New configurations are created by choosing top-ranked configurations from the pool and combining them by means of a genetic engineering procedure as in GGA++, just with a different surrogate and including mutation after the crossover procedure.

## 8. Instance-specific Methods

Offline configurators like ParamILS, irace, GGA, or SMAC employ a one-size-fits-all paradigm, and while this works well for homogeneous problem instance sets where a single configuration yields good performance for all problem instances, it may fall short when the instance set is heterogeneous. For such situations, instance-specific configurators have been developed that can provide a configuration  $\hat{\theta}_i$  for a specific problem instance  $i \sim \mathcal{P}$ .

Instance specific configurators require the same inputs as offline methods with the addition of a feature vector  $\mathbf{f}_i$  for each problem instance (see Figure 5). By altering Equation (2), the problem of instance specific configuration can formally be expressed as:

$$\hat{\theta}^* \in \arg \min_{\theta \in \Theta} \mathbb{E}[c(i, \theta)] \quad , \quad (10)$$

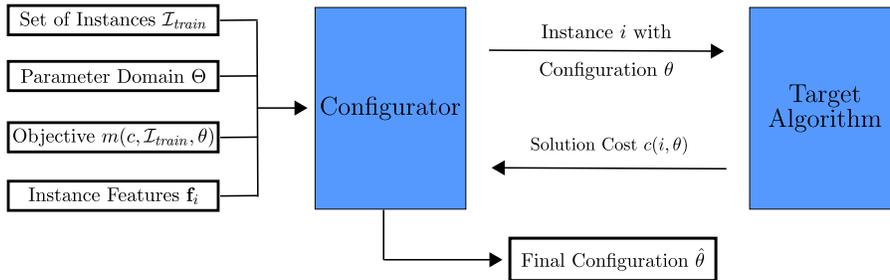


Figure 5: Illustration of Instance Specific Configuration

where  $\mathbb{E}[c(i, \theta)]$  is the expected cost of  $\mathcal{A}$  on a specific problem instance  $i$  for which a configuration should be found. Configurators achieve this goal by harnessing the instance features and learning a relationship between the structure of the instances as described by the features and the performance of various configurations. Thus, the features must provide a good enough representation of the instances such that a model can be constructed, and they need to be correlated with the performance of the target algorithm. This is a difficult task that has itself been the focus of various work (Kroer & Malitsky, 2011; Tierney & Malitsky, 2015).

There are two main ways of generating instance-specific configurations: (1) incorporate the problem instance features into the existing prediction models of the configurator for configurations or (2) deploy separate models that only consider the problem instance features without the configuration characteristics. Both variants can be referred to as model-based, but this section focuses primarily on methods that do not utilize a surrogate model. In Section 5 and 7 model-based and realtime configurators that are instance-specific and use configurations characteristics besides problem instance features in surrogate models are discussed.

**ISAC** Instance-specific algorithm configuration (ISAC) (Kadioglu et al., 2010) and its extensions pair an algorithm configurator with a clustering algorithm. The clustering algorithm partitions the instance space, and the configurator (usually GGA/GGA++) is applied to each of the clusters to generate a configuration specific to the cluster. In the test phase, the learned configurations are assigned to the instances needing to be solved. The assignment mechanism varies depending on the version of ISAC that is used, and is described below in more detail.

ISAC uses the *g-means* (Hamerly & Elkan, 2004) clustering algorithm with a minimum cluster size to generate clusters with the Euclidean distance as a metric. Note that *g-means* does not require the number of clusters to be specified in advance as in other clustering algorithms. Instead, it assumes that clusters ought to be normally distributed and splits clusters that are not using *k-means* with  $k = 2$ . Once the clusters are determined, GGA (Ansótegui et al., 2009) is used to find a representative configuration for each cluster identified by *g-means*, although any configurator could be used. In addition, an extra, default configuration is determined for all instances as a fallback mechanism for instances in the test phase that are not close to the previously determined clusters. In the testing phase,

when a new problem instance arrives, ISAC computes its proximity to each cluster based on the Euclidean distance and uses the configuration of the closest cluster or the default configuration to solve the problem instance.

Malitsky et al. (2013a) introduce evolving instance-specific algorithm configuration (EISAC), which is a retraining mechanism for the test phase that allows for reassignment of configurations to clusters as problem instances arrive. EISAC performs a re-clustering utilizing newly seen instances in addition to the previous training instances, where the update is performed based on the Rand index (Rand, 1971) or when new configurations (solvers) are added/removed. Clusters are updated through solving an optimization problem that redistributes the configurations among clusters by minimizing the solving time of problem instances within a cluster. This means no new configurations are needed. However, additional target algorithm runs would still be needed for each cluster/configuration combination. To avoid this and to save further time, EISAC uses an empirical hardness model to predict the runtime of a configuration on new problem instances. Using this model makes the actual target algorithm runs only necessary when the problem instance would be assigned to a different cluster based on the cost prediction.

Malitsky and Sellmann (2012) extend ISAC to AS. ISAC is used to select an algorithm for a problem instance, and also configures the algorithm automatically, on a per cluster basis. A model-based algorithm selector is added by Ansótegui et al. (2016) in a later version, resulting in the ISAC++ method, which uses cost-sensitive hierarchical clustering (CSHC) (Malitsky et al., 2013b) as a selector.

**CluPaTra** Another offline configurator that, similar to ISAC, derives configurations for clusters of problem instances is CluPaTra (Lau & Lo, 2011; Lindawati et al., 2013a). It is specifically designed for target algorithms that provide search trajectories. It utilizes AGNES (Kaufman & Rousseeuw, 2009) for clustering and ParamILS as a configurator. Instead of using precomputed problem instance features like ISAC, CluPaTra uses the search trajectories derived through runs of the target algorithm encoded as directed sequences. The search trajectory tracks the history of solutions found by the target algorithm through the search space, meaning CluPaTra pierces the black box assumed by most other configurators. A sequence alignment method computes the similarity between trajectories for the clustering. A major drawback of CluPaTra, and its extension, is that it is not fully clear how to utilize it in production. Determining the cluster membership of a new instance is a chicken-and-egg problem, as one must solve the instance to determine the search trajectory that is used for clustering.

The SufTra and FloTra methods extend CluPaTra. The main focus for these approaches lies in different ways of integrating the search trajectories into the clustering. SufTra (Lindawati et al., 2013b) replaces the directed sequence encoding and sequence alignment with a suffix tree encoding from which features are extracted through frequent substring alignment. The similarity between trajectories is computed using frequent substrings and the cosine similarity. In addition, easy problem instances that are similar to hard problem instances are used as surrogates during configuration by finding configurations on easy instances and using them on similar hard instances. This mechanism, however, again struggles from the previously mentioned chicken-and-egg problem, since the runtime needs to be known before being able to partition instances into easy and hard ones. In FloTra (Lindawati et al.,

2013c), the use of the search trajectories is altered by creating graphs out of the trajectories, to which pattern mining is applied to derive features that are used for clustering.

**Hydra** Xu et al. (2010) and Xu et al. (2011) introduce Hydra, which essentially adapts the boosting paradigm (Schapire, 2003) to algorithm configuration. Hydra iteratively configures a solver and adds the configuration to a portfolio that improves upon the weakness of the current portfolio. Hydra utilizes ParamILS (Hutter et al., 2009b) to propose configurations to SATzilla (Xu et al., 2008) which in turn constructs the portfolio. In general, any configurator and any AS technique can be used. The cost function of ParamILS scores configurations by their real value if they perform better than the current portfolio, and by the costs of the portfolio if not. This adjustment effectively penalizes configurations that yield a bad objective value on problem instances the portfolio also performs badly, while not penalizing bad objective values on problem instances on which the portfolio already performs well on. Hydra excels at specifically targeting areas of the feature space in need of custom configurations. However, this comes at a high computational cost as the configurations must be computed sequentially, unlike in ISAC, where they can be parallelized.

**MATE** Genetic programming forms the basis for the model-based algorithm tuning engine (MATE) (Yafrani et al., 2020). The approach takes problem instance features into account and provides human-understandable relations between features and target algorithm parameters. In particular, MATE uses symbolic regression (Augusto & Barbosa, 2000) and tree-based genetic programming. Configurations are encoded as trees, which are compared to each other using a score function that is used to aggregate results over algorithm runs. New trees are generated using genetic operators and replace old trees based on a Wilcoxon test using the measured score and tree complexity. According to the authors, this method should be understood as proof of concept and has not yet been used to configure target algorithms with more than one parameter.

**PCIT** Similar to ISAC, parallel configuration with instance transfer (PCIT) (Liu et al., 2019) partitions problem instances into clusters and finds a configuration for each cluster. While ISAC produces clusters once, and only afterwards finds a configuration for each cluster, PCIT adjusts clusters and assigns configurations sequentially. This allows PCIT to adjust its clusters according to the configurations for each cluster. This comes at the expense of additional configuration costs.

PCIT employs an instance transfer mechanism to shift around instances between groups as new runtime data becomes available. In particular, this mechanism is used in case it becomes clear that a configurator can not find a configuration that is suitable for all instances within a cluster. The transfer mechanism of PCIT has to (i) identify the instances to transfer and (ii) choose a cluster to transfer the instances to. Instances are selected to be transferred when their runtime using the configuration in the respective cluster is higher than the median runtime of all instances over all clusters. If the instance is to be transferred, a surrogate is used to predict the runtime of the instance for the configurations of the other clusters. An instance is assigned to the cluster with the lowest predicted runtime. After the transfer, a configurator is run for every cluster. This loop continues until the computational budget is exhausted. The algorithm configurator SMAC is used to find configurations for each cluster, however any offline configurator can be used.

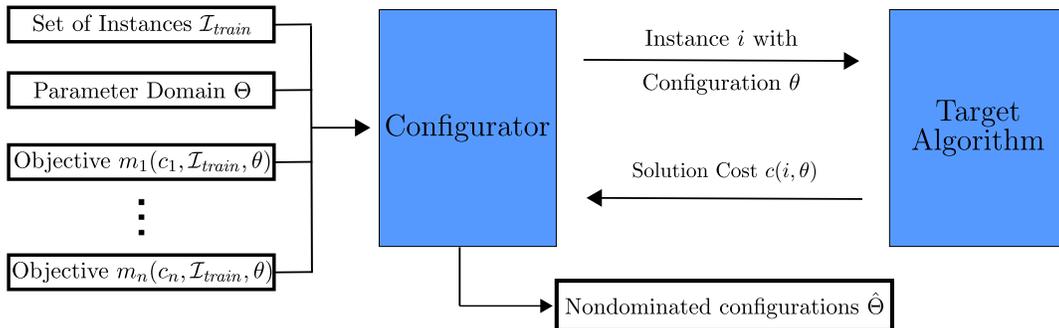


Figure 6: Multi-objective AC

## 9. Multi-objective Methods

We now extend our survey to AC methods that can tune multiple objectives at the same time, such as runtime and quality, or runtime and memory usage (Thanh & Causmaecker, 2014). That is, instead of a single objective,  $m(c, \mathcal{I}_{train}, \theta)$ , multiple objectives  $\mathcal{M} := (m_1, \dots, m_n)$  may be present and competing with each other. The goal of the configurator is then to find a set of configurations  $\hat{\Theta} \subseteq \Theta$  such that no  $\theta \in \hat{\Theta}$  is dominated by another  $\theta'$  with respect to some dominance relation  $\prec$  over the objective functions (Blot et al., 2016). If given preferences from the decision maker, the multiple objectives can be weighted and a single “best” configuration  $\hat{\theta} \subseteq \hat{\Theta}$  can be provided to the decision maker. Note that this AC setting should not be confused with the task of automatically configuring multi-objective target algorithms. The key difference is that the former is concerned with the objective function of the configurator while the latter is interested in the objective function of the target algorithm (Bezerra et al., 2020). Consider that, for example, essentially any single-objective metaheuristic could be configured in a multi-objective way, as quality and runtime present a tradeoff in most of these approaches. The adjustments necessary for single objective configurators to handle multiple configurations are not trivial.

Nevertheless, the multi-objective AC scenario is clearly linked to multi-objective optimization (Coello Coello, 2006). Multi-objective target algorithms can either be configured by means of multi-objective configurators or by choosing a single objective target metric to be optimized by the configurator, such as the hypervolume or epsilon measure (Bezerra et al., 2012; Lopez-Ibanez & Stutzle, 2012; Blot et al., 2017). In the first case, it is not fully clear how the objective values obtained by running the target algorithm translate into configurator objectives. That is, multi-objective target algorithms usually return a set of solutions that produce different values for the different objectives, which in turn would have to be considered by the configurator.

**MO-ParamILS** MO-ParamILS (Blot et al., 2016) is an offline configurator that considers multiple objectives and returns a set of non-dominated configurations for the target algorithm. It is based on the ParamILS framework and, like its predecessor, provides two

evaluation techniques while also utilizing the one exchange neighborhood search. In addition, it can run in parallel and can handle large discrete search spaces.

To handle multiple objectives, MO-ParamILS introduces an archive of configurations and utilizes Pareto dominance to compare configurations. More precisely, it replaces the single configuration that is used during the local search and as incumbent through an archive of non-dominated configurations. In addition, the evaluation mechanism to compare configurations is adjusted to operate on Pareto dominance. In particular, the notion of dominance is used to compare two configurations during local search and also to discard configurations from the archive as new configurations are added.

**Multi-objective racing** Racing procedures can be adjusted to take multiple objectives into account. S-Race (Zhang et al., 2013, 2015a) and SPRINT-Race (Zhang et al., 2015b) are multi-objective configurators similar to F-Race (Birattari et al., 2002; Birattari, 2009), which utilize races and statistical testing to find a set of Pareto optimal configurations for discrete parameter spaces.

S-Race and SPRINT-Race differ from F-Race in terms of the test they apply to discard configurations. Both start with a full factorial design, from which they iteratively eliminate configurations based on the race results. S-Race uses a sign test (Wackerly et al., 2014) to perform a pairwise comparison and eliminate configurations, and Holm’s step-down procedure (Holm, 1979) is applied to account for errors. Utilizing the sign test can lead S-Race to unnecessarily often race and compare two configurations between which no dominance relation exists. To save evaluation time, SPRINT-Race (Zhang et al., 2015b) replaces the sign test with a sequential probability ratio test with indifference zones (Wald, 1947), enabling it to stop running new races when already enough evidence is gathered that a dominance relation is unlikely. Like F-Race, both methods face the drawback of only being able to handle small, discrete configuration spaces. In addition, studies on the competitiveness of the described approaches are limited in scope.

## 10. Dynamic Methods

Dynamic algorithm configuration (Biedenkapp et al., 2020) (DAC) adjusts the configuration *at runtime* instead of committing to a single configuration for solving an entire problem instance. Thus, rather than recommending a single configuration, DAC approaches generate a policy to adapt the configuration dynamically. Note that even realtime AC commits to a single configuration when running a given instance, while DAC has the freedom to adjust the configuration according to target algorithm behavior during execution. Similar to offline AC, DAC can either focus on finding a policy for a set of problem instances or a policy that is tailored towards a single problem instance (i.e., per-instance algorithm configuration).

Two requirements must be met to implement DAC: (1) the algorithm in question needs to support dynamic changes in its configuration and (2) runtime information must be provided to describe the current state of the target algorithm.

DAC approaches consider two different types of features: instance features  $\mathcal{I}$ , which do not change during target algorithm execution, and features encoding the internal state  $\mathcal{Q}$  of the algorithm. Examples of state features include the current iteration of a local search algorithm, the current restart number of a SAT method, or the current solution quality for optimization techniques.

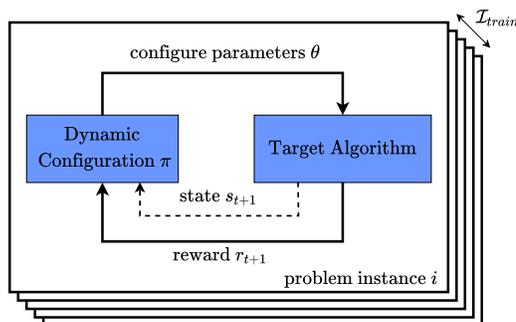


Figure 7: Illustration of dynamic AC as presented in (Biedenkapp et al., 2020)

Biedenkapp et al. (2020) provide the first formal definition of the DAC setting, however, there is a significant amount of earlier work for learning dynamic configuration policies (Lagoudakis & Littman, 2000, 2001; Pettinger & Everson, 2002). Such earlier works use the labels parameter control (Karafotias et al., 2015), online algorithm selection (Vermetten et al., 2019), adaptive selection/configuration (Fialho et al., 2010; van Rijn et al., 2018), Self-Adaptive Monte Carlo Tree Search (Sironi et al., 2018) or hyper-reactive search (Ansótegui et al., 2017, 2018a). For a comprehensive overview of parameter control with respect to parameters of evolutionary algorithms, we refer the interested reader to Karafotias et al. (2015).

van Rijn et al. (2018) and Vermetten et al. (2019) identify potential performance gains through DAC versus static configurations. Based on this insight and data, it is demonstrated in Ye et al. (2021) that performance gains can already be achieved when the algorithm configuration is adapted only once. Furthermore, the hyper-reactive approach of Ansótegui et al. (2017) won several categories at the MaxSAT Evaluation 2016 (Argelich et al., 2016). Thus, DAC offers significant potential for improving algorithms, however, it does require algorithm designers to more deeply integrate their techniques with AC methods than was performed in the past. In the following, we discuss the most frequently used approach to DAC, reinforcement learning (RL). While it is the most popular choice, there also exist other approaches such as policy portfolios, autoconstructive evolution, and multi-armed bandits.

**Reinforcement learning** A challenge in DAC is to evaluate the quality of individual parameter adjustments since feedback is obtained only at the end of processing a problem instance. Due to this, it is natural to consider DAC within the realm of reinforcement learning (we refer to this as DAC-RL). To this end, the DAC problem is modeled in terms of a Markov decision process (MDP) (Puterman, 1994) based on which a policy  $\pi : \mathcal{I} \times \mathcal{Q} \rightarrow \Theta$  is learned from data (Biedenkapp et al., 2020). As a reward signal, typically, the eventually obtained performance value is propagated back to the agent. Using RL allows a wide variety of researched mechanisms to be directly applied to DAC.

The roots of DAC extend to a variety of methods that use RL to control parameters of optimization approaches online. For example, in genetic algorithms (Sakurai et al., 2010; Karafotias et al., 2014), planning algorithms (Pageau et al., 2019; Speck et al., 2021; Bhatia et al., 2021), hyper-heuristics (Özcan et al., 2010), physics simulations (Armstrong et al.,

2006), and evolutionary strategies (Shala et al., 2020). The work of Sakurai et al. (2010), Karafotias et al. (2014) and Özcan et al. (2010) in particular can be widely applied to optimization problems, whereas other works are more application focused.

**Hyper-configurable reactive search** Under the label hyper-configurable reactive search (HCRS), Ansótegui et al. (2017), Ansótegui et al. (2018a) and Sellmann and Tierney (2020) propose an approach to DAC by using standard, offline AC techniques to train a dynamic policy. More specifically, instead of tuning the parameters of an algorithm directly, each parameter is determined by a logistic regression that accepts runtime features from the target algorithm. The parameters of the logistic regressions are exposed to the offline AC technique, GGA++. The target algorithm can thus cheaply query the values it needs, when it needs them, according to the current search state. In Ansótegui et al. (2017), the problem instances are further grouped via the CSHC algorithm (Malitsky et al., 2013b) and for each of the resulting clusters, such a policy encoded via logistic regression models is configured according to the ISAC method (Kadioglu et al., 2010).

**Autoconstructive evolution** In autoconstructive evolution, the aim is to not only evolve a solution to an optimization problem, but also to evolve the genetic operators that are used for recombination, selection, and diversification (Spector & Robinson, 2002; Harrington et al., 2012; Spector & Moscovic, 2017). To this end, each individual represents a tuple of a solution encoding and the encoding of the programs for the genetic operators, respectively. Typically, genetic programming is applied to this problem setting since it can represent programs naturally in terms of abstract syntax trees. In contrast to other algorithm configuration approaches, considering only Boolean, numeric, or nominal parameters of algorithms, the configuration of such genetic operators is even more challenging. Additionally, autoconstructive evolution goes beyond what is usually done in DAC. It adapts the genetic operators for each individual of the population so that the mating behavior changes over time and is thus specific for each individual. Another difference is that for adapting the genetic operators of the evolutionary algorithms, there is no need to engineer and monitor any features representing the algorithm’s state. Although Spector and Moscovic (2017) obtain promising initial results, the combined evolution of solutions and their genetic operators is comparatively challenging, and it remains an open question whether autoconstructive evolution is indeed superior to non-autoconstructive evolution.

**Multi-armed bandits** DAC can be cast as a multi-armed bandit problem, where each configuration represents an arm with an associated reward distribution. The agent is then tasked to maximize the cumulative reward, i.e., “pull the arm“ that maximizes the reward. Put differently, the agent aims to select a configuration that works well for the current time step. This setting of the bandit problem is particularly challenging because the reward distributions associated with the respective arms change over time. In Fialho et al. (2010), various multi-armed bandit approaches are compared with an approach that only considers rewards obtained within a certain time frame (i.e., a sliding window) to allow for smoother transitions between configurations. Note that in this work, the authors assume rewards to be retrievable during an algorithm run. While this assumption holds for many optimization approaches, it does not necessarily hold for all of them (e.g., the pre-processing/root node phase of solving mixed-integer programs), and also often does not hold for simulations.

## 11. Research Directions

Many fruitful avenues of research remain in the area of AC. We formulate research directions with three themes in mind. First, we consider what is required for AC techniques to be more widely adopted in industry. Second, we investigate algorithm configuration settings that have not yet been considered. Third, we consider what methodological advances could move the field forward.

### 11.1 Industry Adoption

Compared to AC, HPO has begun to be adopted in both industry and academia (Van der Blom et al., 2021). Furthermore, HPO is the focus of a number of start-ups, such as SigOpt (which was acquired by Intel (Intel Corporation, 2019)), MindsDB, pecan.ai, and dotData, among many more. AC, however, has received comparatively little fanfare or adoption. While some well-known AC tools have public releases, such as SMAC (Lindauer et al., 2021), irace (López-Ibáñez et al., 2016), and GGA (Ansótegui et al., 2009), we are unaware of widespread use of these in industry. The company Optano GmbH released a version of the GGA configurator (OPTANO GmbH, 2021)<sup>10</sup>, making it perhaps the first commercially developed, general-purpose AC software. The mixed-integer programming solvers Gurobi and CPLEX both contain tools to adjust their parameters. However, to the best of our knowledge, these tools are both outperformed by publicly available research tools. Furthermore, we do not know how either of these tools work.

We speculate that there are several reasons these tools are not frequently used in industry. First, the target algorithm must be carefully designed with its parameters exposed through a command-line interface as proposed by the programming by optimization paradigm (Hoos, 2012). Second, in some cases, users may not think they have sufficient data available to properly configure their approach. This is related to the *cold-start problem* in ML, in that users of AC systems may not have gathered enough instances to properly train their parameterized algorithm to solve their problem effectively. While some work has been performed in this direction in terms of automatically generating instances Malitsky et al. (2016); Smith-Miles and Bowly (2015); Akgün et al. (2019); Tang et al. (2021); Liu et al. (2022), there are still many real-world problems that cannot be modeled with these techniques.

Third, the runtimes of some target algorithms are too high for AC. In these cases, transfer learning from a dataset of similar, but easier, problem instances may be a way forward. Alternatively, the AC algorithms could try to guess the quality of the target algorithm before it is finished, or at least guess a ranking over the different configurations.

A fourth reason for a lack of adoption in industry is likely the inability of current algorithm configuration tools to integrate into existing systems. Target algorithms may not be easily accessible over the command line or accept data in a simple “problem instance” format as required by existing configurators. For example, if an algorithm is tightly integrated with a database or SAP system, AC may require significant extra work. The solution here is likely not a new research concept or method, but rather raising awareness that de-

---

10. The configurator from Optano GmbH is the result of a joint research project between Optano GmbH and Kevin Tierney, but we note that it is freely available and there is no commercial interest on the part of Kevin Tierney.

coupling solvers from production environments will enable configuration of the algorithm’s parameters.

A final, fifth reason is the expert knowledge required to set up an AC environment. AC is not a topic that is widely taught in data science or computer science curriculums, although at least HPO is beginning to be adopted. Nonetheless, a deep familiarity with the various components of AC is necessary to successfully implement AC in practice. Furthermore, we are not aware of many consultants offering AC services to fill this knowledge gap.

## 11.2 Novel AC Settings

We identify several AC settings according to our problem view that have not yet been realized in the literature. In particular, settings considering multiple objectives, especially in terms of the target algorithm, during run target algorithm observations, dynamic configuration adjustment, and parameter transfer learning have not been significantly considered in the literature. It is easy to imagine real-world applications for all the situations listed, thus there remain ample research opportunities available in AC.

**Multiple objectives** In Section 9, we identify several approaches for multi-objective AC, but note that these are rather limited in how they incorporate multiple objectives into their configuration. Especially in regard to runtime objectives, multi-objective AC methods are still lacking adaptive mechanisms for capping runs or generating configurations that target specific areas of the Pareto front. While we expect multi-objective target algorithms to integrate into these multi-objective AC techniques, extra heuristics or care could be taken to consider the Pareto front of the target algorithm in addition to (or instead of) a Pareto front over configurations. Furthermore, multiple objectives have only been included in offline configuration.

Since realtime AC is tasked with returning solutions directly to users, providing multiple solutions in a multi-objective context could be valuable. Realtime configurators only get one shot at solving a problem instance (or, in this case, generating multiple solutions). Thus, the research focus would be on generating an interesting set of solutions for the user. This set ought to adequately represent the Pareto front while providing the solutions quickly if runtime is one of the considered objectives.

**Incorporating runtime information & DAC** While adjusting algorithm configurations dynamically is not a new concept, learning policies with AC techniques is, and as pointed out in Section 10, has shown considerable promise. These techniques use information from the target algorithm about its current state to help create an effective policy for setting parameters. Many algorithms offer indications as to how good (or bad) their solution procedure is progressing, especially when compared to the output of other configurations solving the same instance. However, little work has been done on using this information in the AC approach itself, e.g., to stop unpromising configurations early or to assist in comparing configurations with censored runtime data. This line of work has the potential to reduce the overall computation time of AC.

**Meta-AC** AC methods themselves contain many parameters that require tuning, but, in a twist of irony, the computational cost of configuring a configurator makes this a daunting task. The AS problem has been similarly investigated on a meta-level by Tornede

et al. (2020a, 2021b) for learning ensembles of algorithm selectors. Moreover, this sort of “meta configuration” has been performed by Ansótegui et al. (2018b) to configure algorithm selection portfolios with GGA and Lindauer et al. (2017). Given that considering the “hyper”-hyperparameters in HPO has seen also some success (Feurer & Hutter, 2018), we see potential for the general AC setting as well.

**Transfer learning** Consider a target algorithm for which a new version is released with several new parameters, and perhaps some existing parameters have different domains or are removed. From a practical standpoint, there are ways of handling this in modern AC software, such as inserting the previously tuned parameters along with some default values for the rest as a starting point of the search. Of course, this involves starting a brand-new configuration process that “forgets” everything that was previously learned about how the previous parameters behaved on a set of instances. Maintaining this state and transferring knowledge to the new or modified parameters could result in better configurations and less computation time used on AC. A first step towards this was recently proposed by Franzin and Stützle (2020).

**Generalization** Configurations found by offline configurators (see Section 4 and 5) might not always perform as well in production as one might expect based on the achieved training performance. That is, the configurations found may not generalize well to new instances or performance may degrade over time due to concept drift. To address this, the community has shifted focus towards methods that allow for an active control of the configuration during production based on the new instance, e.g., in the realtime and dynamic configuration settings. Nonetheless, more attention is needed to ensuring that configurations remain effective in production settings.

### 11.3 Novel Benchmarks

To accelerate future research and to provide reproducible results, benchmarks for new problem settings are needed. AClib (Hutter et al., 2014a), the currently most widely used benchmark library in the context of AC, covers the offline configuration case and provides a set of different problems (SAT, MIP, ASP, etc.) of varying complexity (number of variables and problem instances) for tasks of runtime or quality configuration. For DAC, the DACBench has been proposed (Eimer et al., 2021), although this does not support DAC settings envisioned, e.g., by hyper-reactive search. As an alternative to such libraries, AC methods can also be benchmarked by using surrogate models that are trained on test instances in advance, resulting in cheaper evaluations when testing (Eggenesperger et al., 2018). The existing benchmarks fail to cover other configurations settings like the realtime configuration setting or the configuration of multi-objective target algorithms.

### 11.4 Novel AC Methodologies

AC methods have become extremely sophisticated and cover a wide range of methodologies including evolutionary algorithms, statistical tests, and learned surrogate models. There nonetheless remain opportunities to improve current methods and create better AC algorithms. We note that our goal in this section is not necessarily to specify the methodologies of the future, but rather to identify the key challenges that remain in the hopes that engaged

readers will fill these gaps with their ideas. To this end, we discuss several challenges faced by AC approaches: comparing/ranking configurations, generating new configurations, and selecting instances.

**Comparing/ranking configurations** This challenge can be summarized as follows: given two or more configurations, determine which one(s) is(are) the best performing without needing to run the configurations on the entire training set of instances. In offline configuration, the ranking can be used to generate new configurations (part of the next challenge), whereas in realtime configuration the rankings can be critical to deciding which configurations are allowed to try to solve the current instance. Methods for ranking and comparing include using empirical hardness models Leyton-Brown et al. (2009) or more general surrogates as in GGA++ (Ansótegui et al., 2015) or SMAC (Hutter et al., 2011), statistics (López-Ibáñez et al., 2016), the TrueSkill mechanism (Fitzgerald et al., 2015), bandits (El Mesaoudi-Paul et al., 2020b). Nonetheless, there is undoubtedly still room for improvement, using perhaps new preference learning techniques (as have been used for AS in Hanselle et al. (2020)) or deep learning models.

**Generating new configurations** Every AC algorithm must have a mechanism for generating new configurations. The question, of course, is how to generate configurations that are high quality with respect to the configuration objective. This research question is closely related to comparing and ranking configurations, as mechanisms like the one in GGA++ compare configurations that are generated according to some rules. Creating new configurations, especially in an instance-specific capacity, offers a clear path to higher quality AC mechanisms that can be used across a wide range of AC problem settings.

**Instance selection** This review begins by identifying instance handling mechanisms as one of the key differences between HPO and general AC. The strategies used in the literature are fairly basic. Most algorithms either select random subsets of the instances or expand a subset. It might be worthwhile to adapt methods related to the concept of dataset distillation (Wang et al., 2018) to generate small instance sets that are nonetheless representative of the complete instance set to speed up configuration. Instance-specific approaches usually find some way to partition the instances before tuning multiple configurations. However, much remains unknown about how certain orders of considering instances could change or improve the configurations found. It is possible that the order in which instances are examined does not matter, but given that this is a core competency of general AC methods, the potential for performance improvements is present.

**Surrogate model features** Learning suitable features for surrogate models automatically is closely related to instance selection. Even though there are good features for particular problem classes such as SAT or MILP problems Hutter et al. (2014b), other not so well studied problems lack well-defined features and require domain experts to manually derive them for a specific use case. Previous works (Kroer & Malitsky, 2011; Tierney & Malitsky, 2015) have begun to address this problem in the realm of AS. Adopting this for AC, however, remains an open challenge.

**Bounded rationality and rational metareasoning** Hüllermeier et al. (2021) elaborate on automated machine learning (as well as related problems of automated algorithm

design, such as hyperparameter optimization and algorithm selection) from the perspective of bounded rationality. The authors propose to view methods for automated algorithm design as intelligent agents, which need to take decisions under bounded resources (e.g., configuration time in the case of AC) and thus apply reasoning on a meta-level to decide how to optimally allocate the available resources. The motivation for adopting this perspective is twofold: first, to shed light on existing AutoML methods, and second, to inspire new approaches based on established (decision-theoretic) principles of rational metareasoning and bounded optimality (Russell & Wefald, 1991; Russell, 1997; Zilberstein, 2011; Russell, 2016).

### 11.5 Theoretical Results

Although the field of AC has seen significant progress recently in terms of answering theoretical questions, there are still a number of critical challenges that remain. In the following, we address some of these challenges.

**Generalization guarantees** The recent work by Pushak and Hoos (2018) shows empirically that the configuration landscape, i.e., the class of cost functions, exhibits an approximately uni-modal or convex structure for a couple of NP-hard problem classes and sophisticated target algorithms. In light of this, it seems reasonable to extend the current results regarding generalization guarantees to the scenario, where the functions in the class of cost functions are (approximately) uni-modal or convex. Such an assumption will likely lead to better theoretical generalization guarantees as the current ones reviewed in Section 6.1, as uni-modalities and/or convexity of target functions are known to accelerate the learning process in general.

**Theoretical analysis of state-of-the-art configurators** Although the work of Hall et al. (2019) provides valuable insights into the theoretical properties of ParamILS, it still does not cover all its theoretical properties. Even less is known about the theoretical properties of other state-of-the-art algorithm configurators, such as SMAC, irace, or GGA, besides that SMAC will find the optimal configuration if the configuration space is finite and sufficient time for running SMAC is available (see Theorem 4 in Hutter et al. (2010a)). A deeper theoretical analysis of these algorithms is likely to be helpful to understand the reasons for the practically appealing behavior of these configurators. Furthermore, this analysis could reveal on which classes of problems the configurators are likely to perform effectively, and on which they are not.

**Instance specific theoretical analyses** So far, the existing theoretical approaches do not take additional side information in the form of feature vectors of the problem instances into account, i.e., they all employ a one-fits-all paradigm. However, as several of the reviewed works in Section 8 have shown, there might be a benefit regarding the performance of an algorithm configurator in practical applications if these features are incorporated. In light of this, it would be interesting to investigate theoretically the potential improvement of an instance-specific (i.e., feature-based) configurator over a one-fits-all paradigm (i.e., feature-free) configurator for specific problem scenarios.

## 12. Conclusion

Parameters are ubiquitous in modern optimization approaches and beyond, with all of the significant solvers for, e.g., MILP, SAT, or TSP problems containing parameters that influence their performance and need to be set by the user. AC frees the user from this tedious and error-prone task by automating the search for high-quality configurations. This survey presented an overview of the current state of AC research by outlining relevant methods, their design features, and problem settings. In particular, we provided two taxonomies for organizing AC approaches, reviewed and contrasted different approaches in the light of the problem setting they were designed for, and highlight underlying principles and ideas. Overall, we find that the AC literature is in a mature state, including powerful empirical approaches available to handle a variety of large-scale, real-world challenges, as well as theoretical approaches providing quality guarantees and an understanding of the AC domain.

We identify that the methodological trend is towards incorporating learned models into AC methods and that this has improved performance in every AC setting where it has been attempted. Nonetheless, the AC literature shows a surprising amount of hybridization of local search, evolutionary and model-based methods. We hypothesize that there is still significant progress that can be made in the area of AC, despite the sophistication of current methods, and are encouraged by the significant increase in attention the field has received, in particular through the spread of HPO techniques. Finally, we especially encourage researchers to address the real-world usability of AC techniques to ensure that the promising performance gains the AC community is seeing can benefit the world at large.

## Acknowledgements

This work was partially supported by the German Research Foundation (DFG) within the Collaborative Research Center “On-The-Fly Computing” (SFB 901/3 project no. 160364472) and by the research training group “Dateninja” (Trustworthy AI for Seamless Problem Solving: Next Generation Intelligence Joins Robust Data Analysis) funded by the German federal state of North Rhine-Westphalia. We also would like to thank the anonymous reviewers for their suggestions on the first version of this manuscript.

## 13. Appendix

To help the reader navigate though the jungle of AC, we provide additional resources. Table 6 contains a list of abbreviations with terms related to AC used within this work. In addition, we provide a list of software resources (Table 7) that contains currently available tools for AC. We only include software that is widely used.

<b>General</b>	
AC	Algorithm configuration
AS	Algorithm selection
BN	Bayesian networks
BO	Bayesian optimization
CASH	Combined algorithm selection and hyperparameter optimization
DAC	Model-based algorithm tuning engine
DOE	Design of experiments methodology
EA	Evolutionary algorithm
EI	Expected improvement
GP	Gaussian processe
HPO	Hyperparameter optimization
HCRS	Hyper-configurable reactive search
ILS	Iterated local search
MILP	Mixed-integer programming
RVNS	Reduced variable neighborhood search
SAT	Boolean satisfiability problems
SMBO	Sequential model-based optimization
UCB	Upper confidence bound
<b>Method related</b>	
BNT	Bayesian network tuning
CAR	CapsAndRuns
CPPL	Contextual preselection with Plackett-Luce
D-SMAC	Distributed SMAC
EISAC	Evolving instance-specific algorithm configuration
GGA	Gender-based genetic algorithm
GPS	Golden Parameter Search
HORA	Heuristic oriented racing algorithm
ICAR	ImpatientCapsAndRuns
I/F-Race	Iterated F-Race
ISAC	Instance-specific algorithm configuration
MATE	Model-based algorithm tuning engine
MBGM	Model-based graphical methods
ReACT	Realtime algorithm configuration through tournaments
REVAC	Parameter relevance estimation and value calibration
ROAR	Random online aggressive racing
SKO	sequential kriging meta-modelling
SMAC+PS	Sequential model-based optimization for algorithm configuration
SMAC	SMAC and probabilistic sampling
SPO	Sequential parameter optimization
SP	Structured procrastination
TB-SPO	Time-bounded SPO

Table 6: Glossary of acronyms used in this work

<b>General AC systems</b>	
D-SMAC	<a href="https://github.com/tqichun/distributed-SMAC3">https://github.com/tqichun/distributed-SMAC3</a>
GPS	<a href="https://github.com/YashaPushak/GPS">https://github.com/YashaPushak/GPS</a>
irace	<a href="https://github.com/MLopez-Ibanez/irace">https://github.com/MLopez-Ibanez/irace</a>
OAT (GGA)	<a href="https://docs.optano.com/algorithm.tuner/current/">https://docs.optano.com/algorithm.tuner/current/</a>
ParamILS	<a href="https://www.cs.ubc.ca/labs/algorithms/Projects/ParamILS/">https://www.cs.ubc.ca/labs/algorithms/Projects/ParamILS/</a>
PyDGGGA	<a href="http://ulog.udl.cat/software/">http://ulog.udl.cat/software/</a>
REVAC	<a href="https://github.com/ChrisTimperley/RubyREVAC">https://github.com/ChrisTimperley/RubyREVAC</a>
SMAC 3	<a href="https://github.com/automl/SMAC3">https://github.com/automl/SMAC3</a>
<b>Benchmarks</b>	
AClib	<a href="https://bitbucket.org/mlindauer/aclib2/src/master/">https://bitbucket.org/mlindauer/aclib2/src/master/</a>
DAC	<a href="https://github.com/automl/DAC">https://github.com/automl/DAC</a>

Table 7: List of available software in the realm of AC.

## References

- Belarmino Adenso-Diaz & Manuel Laguna. Fine-tuning of algorithms using fractional experimental designs and local search. *Operations research*, 54(1):99–114, 2006.
- Özgür Akgün, Nguyen Dang, Ian Miguel, András Z Salamon, & Christopher Stone. Instance generation via generator instances. In *International Conference on Principles and Practice of Constraint Programming*, pages 3–19. Springer, 2019.
- Marie Anastacio & Holger H. Hoos. Model-based algorithm configuration with default-guided probabilistic sampling. In *International Conference on Parallel Problem Solving from Nature*, pages 95–110. Springer, 2020.
- Carlos Ansótegui, Meinolf Sellmann, & Kevin Tierney. A gender-based genetic algorithm for the automatic configuration of algorithms. In *International Conference on Principles and Practice of Constraint Programming*, pages 142–157. Springer, 2009.
- Carlos Ansótegui, Yuri Malitsky, Horst Samulowitz, Meinolf Sellmann, & Kevin Tierney. Model-based genetic algorithms for algorithm configuration. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI*, pages 733–739, 2015.
- Carlos Ansótegui, Joel Gabas, Yuri Malitsky, & Meinolf Sellmann. MaxSAT by improved instance-specific algorithm configuration. *Artificial Intelligence*, 235:26–39, 2016.
- Carlos Ansótegui, Josep Pon, Meinolf Sellmann, & Kevin Tierney. Reactive dialectic search portfolios for MaxSAT. In *Proceedings of the Thirty-First Conference on Artificial Intelligence, AAAI*, pages 765–772. AAAI Press, 2017.
- Carlos Ansótegui, Britta Heymann, Josep Pon, Meinolf Sellmann, & Kevin Tierney. Hyper-reactive tabu search for MaxSAT. In *Learning and Intelligent Optimization - 12th International Conference, LION*, volume 11353 of *Lecture Notes in Computer Science*, pages 309–325. Springer, 2018a.

- Carlos Ansótegui, Meinolf Sellmann, & Kevin Tierney. Self-configuring cost-sensitive hierarchical clustering with recourse. In *Principles and Practice of Constraint Programming*, pages 524–534. Springer, 2018b.
- Carlos Ansótegui, Josep Pon, & Meinolf Sellmann. Boosting evolutionary algorithm configuration. *Annals of Mathematics and Artificial Intelligence*, pages 1–20, 2021.
- Carlos Ansótegui, Josep Pon, Meinolf Sellmann, & Kevin Tierney. PyDGGA: Distributed GGA for automatic configuration. In *Theory and Applications of Satisfiability Testing - SAT*, volume 12831 of *Lecture Notes in Computer Science*, pages 11–20. Springer, 2021.
- Josep Argelich, Chu Min Li, Felip Manyà, & Jordi Planes. MaxSAT Evaluation 2016, 2016. URL <http://maxsat.ia.udl.cat/introduction/>.
- Warren Armstrong, Peter Christen, Eric McCreath, & Alistair P. Rendell. Dynamic algorithm selection using reinforcement learning. In *International Workshop on Integrating AI and Data Mining*, pages 18–25, 2006.
- Gilles Audemard & Laurent Simon. Predicting learnt clauses quality in modern SAT solvers. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI*, pages 399–404, 2009.
- Gilles Audemard & Laurent Simon. Refining restarts strategies for SAT and UNSAT. In *International Conference on Principles and Practice of Constraint Programming*, pages 118–126. Springer, 2012.
- Douglas Adriano Augusto & Helio J.C. Barbosa. Symbolic regression via genetic programming. In *Proceedings of the Sixth Brazilian Symposium on Neural Networks*, pages 173–178. IEEE, 2000.
- Prasanna Balaprakash, Mauro Birattari, & Thomas Stützle. Improvement strategies for the F-Race algorithm: Sampling design and iterative refinement. In *International workshop on hybrid metaheuristics*, pages 108–122. Springer, 2007.
- Maria-Florina Balcan, Vaishnavh Nagarajan, Ellen Vitercik, & Colin White. Learning-theoretic foundations of algorithm configuration for combinatorial partitioning problems. In *Conference on Learning Theory*, pages 213–274. PMLR, 2017.
- Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, & Ellen Vitercik. Learning to branch. In *Proceedings of the 35th International Conference on Machine Learning, ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 353–362. PMLR, 2018.
- Maria-Florina Balcan, Dan F. DeBlasio, Travis Dick, Carl Kingsford, Tuomas Sandholm, & Ellen Vitercik. How much data is sufficient to learn high-performing algorithms? *CoRR*, abs/1908.02894, 2019. URL <http://arxiv.org/abs/1908.02894>.
- Maria-Florina Balcan, Travis Dick, & Manuel Lang. Learning to link. In *8th International Conference on Learning Representations, ICLR*. OpenReview.net, 2020a.

- Maria-Florina Balcan, Tuomas Sandholm, & Ellen Vitercik. Learning to optimize computational resources: Frugal training with generalization guarantees. In *The Thirty-Fourth Conference on Artificial Intelligence, AAAI*, pages 3227–3234. AAAI Press, 2020b.
- Maria-Florina Balcan, Tuomas Sandholm, & Ellen Vitercik. Refined bounds for algorithm configuration: The knife-edge of dual class approximability. In *Proceedings of the 37th International Conference on Machine Learning, ICML*, volume 119, pages 580–590. PMLR, 2020c.
- Adrian Balint & Norbert Manthey. Sparrowtoriss.s. *Proceedings of SAT Competition 2014*, pages 87–88, 2013.
- Eduardo B. M. Barbosa & Edson Luiz França Senne. A heuristic for optimization of metaheuristics by means of statistical methods. In *Proceedings of the 6th International Conference on Operations Research and Enterprise Systems, ICORES*, pages 203–210. SciTePress, 2017a.
- Eduardo B.M. Barbosa & Edson Luiz França Senne. Improving the fine-tuning of metaheuristics: an approach combining design of experiments and racing algorithms. *Journal of Optimization*, 2017, 2017b.
- Peter L. Bartlett & Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.*, 3(null):463–482, mar 2003. ISSN 1532-4435.
- Thomas Bartz-Beielstein. *Experimental analysis of evolution strategies: Overview and comprehensive introduction*. Citeseer, 2003.
- Thomas Bartz-Beielstein, Christian W.G. Lasarczyk, & Mike Preuß. Sequential parameter optimization. In *IEEE congress on evolutionary computation*, volume 1, pages 773–780. IEEE, 2005.
- Viktor Bengs & Eyke Hüllermeier. Preselection bandits. In *Proceedings of the International Conference on Machine Learning, ICML*, pages 778–787, 2020.
- Viktor Bengs, Róbert Busa-Fekete, Adil El Mesaoudi-Paul, & Eyke Hüllermeier. Preference-based online learning with dueling bandits: A survey. *Journal of Machine Learning Research*, 22:1–108, 2021.
- James Bergstra, Rémi Bardenet, Yoshua Bengio, & Balázs Kégl. Algorithms for hyperparameter optimization. In *25th Annual Conference on Neural Information Processing Systems*, pages 2546–2554, 2011.
- Leonardo C. T. Bezerra, Manuel López-Ibáñez, & Thomas Stützle. Automatic generation of multi-objective ACO algorithms for the biobjective knapsack. In Marco Dorigo et al., editors, *Swarm Intelligence, 8th International Conference, ANTS 2012*, volume 7461 of *Lecture Notes in Computer Science*, pages 37–48. Springer, Heidelberg, 2012. doi: 10.1007/978-3-642-32650-9\_4.

- Leonardo C. T. Bezerra, Manuel López-Ibáñez, & Thomas Stützle. Automatic configuration of multi-objective optimizers and multi-objective configuration. In Thomas Bartz-Beielstein, Bogdan Filipič, P. Korošec, & El-Ghazali Talbi, editors, *High-Performance Simulation-Based Optimization*, pages 69–92. Springer International Publishing, Cham, Switzerland, 2020. doi: 10.1007/978-3-030-18764-4\_4.
- Abhinav Bhatia, Justin Svegliato, & Shlomo Zilberstein. Tuning the hyperparameters of anytime planning: A deep reinforcement learning approach. In *Workshop on Heuristics and Search for Domain-independent Planning, ICAPS*, 2021.
- André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, & Marius Lindauer. Dynamic algorithm configuration: Foundation of a new meta-algorithmic framework. In *24th European Conference on Artificial Intelligence, ECAI*, volume 325, pages 427–434. IOS Press, 2020.
- Mauro Birattari. On the estimation of the expected performance of a metaheuristic on a class of instances—how many instances, how many runs? *Tech. Rep. TR/IRIDIA/2004-001*, 2004.
- Mauro Birattari. *Tuning Metaheuristics - A Machine Learning Perspective*, volume 197 of *Studies in Computational Intelligence*. Springer, 2009.
- Mauro Birattari, Thomas Stützle, Luís Paquete, & Klaus Varrentrapp. A racing algorithm for configuring metaheuristics. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO*, pages 11–18. Morgan Kaufmann Publishers, 2002.
- Mauro Birattari, Zhi Yuan, Prasanna Balaprakash, & Thomas Stützle. F-Race and iterated F-Race: An overview. *Experimental methods for the analysis of optimization algorithms*, pages 311–336, 2010.
- Mauro Birattari, Marco Chiarandini, Marco Saerens, & Thomas Stützle. Learning graphical models for parameter tuning. Technical report, Citeseer, 2011.
- Bernd Bischl, Martin Binder, Michel Lang, Tobias Pielok, Jakob Richter, Stefan Coors, Janek Thomas, Theresa Ullmann, Marc Becker, Anne-Laure Boulesteix, Difan Deng, & Marius Lindauer. Hyperparameter optimization: Foundations, algorithms, best practices and open challenges. *CoRR*, abs/2107.05847, 2021.
- Aymeric Blot, Holger H. Hoos, Laetitia Jourdan, Marie-Éléonore Kessaci-Marmion, & Heike Trautmann. MO-ParamILS: A multi-objective automatic algorithm configuration framework. In *International Conference on Learning and Intelligent Optimization*, pages 32–47. Springer, 2016.
- Aymeric Blot, Alexis Pernet, Laetitia Jourdan, Marie-Éléonore Kessaci-Marmion, & Holger H. Hoos. Automatically configuring multi-objective local search using multi-objective optimisation. In *International Conference on Evolutionary Multi-Criterion Optimization*, pages 61–76. Springer, 2017.
- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

- Leslie Pérez Cáceres, Bernd Bischl, & Thomas Stützle. Evaluating random forest models for irace. In *Genetic and Evolutionary Computation Conference*, pages 1146–1153. ACM, 2017.
- Leslie Pérez Cáceres Cáceres & Thomas Stützle. Exploring variable neighborhood search for automatic algorithm configuration. *Electronic Notes in Discrete Mathematics*, 58: 167–174, 2017.
- Carlos A. Coello Coello. Evolutionary multi-objective optimization: a historical view of the field. *IEEE Computational Intelligence Magazine*, 1(1):28–36, 2006.
- Steven P. Coy, Bruce L. Golden, George C. Runger, & Edward A. Wasil. Using experimental design to find effective parameter settings for heuristics. *Journal of Heuristics*, 7(1):77–97, 2001.
- Marcelo De Souza, Marcus Ritt, & Manuel López-Ibáñez. Capping methods for the automatic configuration of optimization algorithms. *Computers & Operations Research*, 139: 105615, 2022. doi: 10.1016/j.cor.2021.105615.
- Marcelo Branco do Nascimento & Antonio Augusto Chaves. An automatic algorithm configuration based on a bayesian network. In *IEEE Congress on Evolutionary Computation, CEC*, pages 1–8. IEEE, 2020.
- Katharina Eggenberger, Marius Lindauer, Holger H. Hoos, Frank Hutter, & Kevin Leyton-Brown. Efficient benchmarking of algorithm configurators via model-based surrogates. *Machine Learning*, 107(1):15–41, 2018.
- Katharina Eggenberger, Marius Lindauer, & Frank Hutter. Pitfalls and best practices in algorithm configuration. *Journal of Artificial Intelligence Research*, 64:861–893, 2019.
- Katharina Eggenberger, Kai Haase, Philipp Müller, Marius Lindauer, & Frank Hutter. Neural model-based optimization with right-censored observations. *arXiv preprint arXiv:2009.13828*, 2020.
- Agoston E. Eiben & Selmar K. Smit. Evolutionary algorithm parameters and methods to tune them. In *Autonomous search*, pages 15–36. Springer, 2011a.
- Agoston E. Eiben & Selmar K. Smit. Parameter tuning for configuring and analyzing evolutionary algorithms. *Swarm and Evolutionary Computation*, 1(1):19–31, 2011b.
- Theresa Eimer, André Biedenkapp, Maximilian Reimer, Steven Adriaensen, Frank Hutter, & Marius Lindauer. DACBench: A benchmark library for dynamic algorithm configuration. *arXiv preprint arXiv:2105.08541*, 2021.
- Adil El Mesaoudi-Paul, Viktor Bengs, & Eyke Hüllermeier. Online preselection with context information under the plackett-luce model. *arXiv preprint arXiv:2002.04275*, 2020a.
- Adil El Mesaoudi-Paul, Dimitri Weiß, Viktor Bengs, Eyke Hüllermeier, & Kevin Tierney. Pool-based realtime algorithm configuration: A preselection bandit approach. In *International Conference on Learning and Intelligent Optimization*, pages 216–232. Springer, 2020b.

- Radwa Elshawi, Mohamed Maher, & Sherif Sakr. Automated machine learning: State-of-the-art and open challenges. *arXiv preprint arXiv:1906.02287*, 2019.
- Yasemin Eryoldaş & Alptekin Durmuşoğlu. A literature survey on instance specific algorithm configuration methods. In *Proceedings of the 11th Annual International Conference on Industrial Engineering and Operations Management*, pages 2983–2990. IEOM Society International, 2021.
- Mehdi Fallahi, Somayeh Amiri, & Masoud Yaghini. A parameter tuning methodology for metaheuristics based on design of experiments. *International Journal of Engineering and Technology Sciences*, 2(6):497–521, 2014.
- Matthias Feurer & Frank Hutter. Towards further automation in AutoML. In *International Conference on Machine Learning, ICML, AutoML workshop*, page 13, 2018.
- Álvaro Fialho, Luís Da Costa, Marc Schoenauer, & Michèle Sebag. Analyzing bandit-based adaptive operator selection mechanisms. *Annals of Mathematics and Artificial Intelligence*, 60(1-2):25–64, 2010.
- Ronald Aylmer Fisher. The design of experiments. *The design of experiments*, 1937.
- Tadhg Fitzgerald, Yuri Malitsky, Barry O’Sullivan, & Kevin Tierney. React: Real-time algorithm configuration through tournaments. In *Proceedings of the Seventh Annual Symposium on Combinatorial Search, SOCS*. AAAI Press, 2014.
- Tadhg Fitzgerald, Yuri Malitsky, & Barry O’Sullivan. Reactr: Realtime algorithm configuration through tournament rankings. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI*, pages 304–310. AAAI Press, 2015.
- Alberto Franzin & Thomas Stützle. Towards transferring algorithm configurations across problems. In *Learning Meets Combinatorial Algorithms at NeurIPS2020*, 2020.
- Alberto Franzin, Leslie Pérez Cáceres Cáceres, & Thomas Stützle. Effect of transformations of numerical parameters in automatic algorithm configuration. *Optimization Letters*, 12(8):1741–1753, 2018.
- Peter I. Frazier. A tutorial on bayesian optimization. *CoRR*, abs/1807.02811, 2018.
- João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, & Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM computing surveys, CSUR*, 46(4):1–37, 2014.
- Aldy Gunawan & Hoong Chuin Lau. Fine-tuning algorithm parameters using the design of experiments approach. In *International Conference on Learning and Intelligent Optimization*, pages 278–292. Springer, 2011.
- Aldy Gunawan, Hoong Chuin Lau, & Elaine Wong. Real-world parameter tuning using factorial design with parameter decomposition. In *Advances in Metaheuristics*, pages 37–59. Springer, 2013.

- George T. Hall, Pietro S. Oliveto, & Dirk Sudholt. On the impact of the cutoff time on the performance of algorithm configurators. *CoRR*, abs/1904.06230, 2019. URL <http://arxiv.org/abs/1904.06230>.
- George T. Hall, Pietro S. Oliveto, & Dirk Sudholt. Analysis of the performance of algorithm configurators for search heuristics with global mutation operators. *CoRR*, abs/2004.04519, 2020a. URL <https://arxiv.org/abs/2004.04519>.
- George T. Hall, Pietro Simone Oliveto, & Dirk Sudholt. Fast perturbative algorithm configurators. *CoRR*, abs/2007.03336, 2020b. URL <https://arxiv.org/abs/2007.03336>.
- Greg Hamerly & Charles Elkan. Learning the k in k-means. *Advances in neural information processing systems*, 16:281–288, 2004.
- Jonas Hanselle, Alexander Tornede, Marcel Wever, & Eyke Hüllermeier. Hybrid ranking and regression for algorithm selection. In *KI 2020: Advances in Artificial Intelligence - 43rd German Conference on AI*. Springer, 2020.
- Jonas Hanselle, Alexander Tornede, Marcel Wever, & Eyke Hüllermeier. Algorithm selection as superset learning: Constructing algorithm selectors from imprecise performance data. In *Advances in Knowledge Discovery and Data Mining - 25th Pacific-Asia Conference, PAKDD*, volume 12712 of *Lecture Notes in Computer Science*, pages 152–163. Springer, 2021.
- Kyle Ira Harrington, Lee Spector, Jordan B. Pollack, & Una-May O’Reilly. Autoconstructive evolution for structural problems. In *Genetic and Evolutionary Computation Conference, GECCO*, pages 75–82. ACM, 2012.
- Kyle Robert Harrison, Beatrice M. Ombuki-Berman, & Andries P. Engelbrecht. The parameter configuration landscape: A case study on particle swarm optimization. In *IEEE Congress on Evolutionary Computation, CEC*, pages 808–814. IEEE, 2019.
- Ralf Herbrich, Tom Minka, & Thore Graepel. Trueskill<sup>tm</sup>: A bayesian skill rating system. In *Advances in Neural Information Processing Systems*, pages 569–576. MIT Press, 2006.
- Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70, 1979.
- Holger H Hoos. Automated algorithm configuration and parameter tuning. In *Autonomous search*, pages 37–71. Springer, 2011.
- Holger H. Hoos. Programming by optimization. *Communications of the ACM*, 55(2):70–80, 2012.
- Changwu Huang, Yuanxiang Li, & Xin Yao. A survey of automatic parameter tuning methods for metaheuristics. *IEEE transactions on evolutionary computation*, 24(2):201–216, 2019.
- Deng Huang, Theodore T. Allen, William I. Notz, & Ning Zeng. Global optimization of stochastic black-box systems via sequential kriging meta-models. *Journal of global optimization*, 34(3):441–466, 2006.

- Jinbo Huang. The effect of restarts on the efficiency of clause learning. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI*, pages 2318–2323, 2007.
- Eyke Hüllermeier, Felix Mohr, Alexander Tornede, & Marcel Wever. Automated machine learning, bounded rationality, and rational metareasoning. *CoRR*, abs/2109.04744, 2021.
- Frank Hutter & Youssef Hamadi. Parameter adjustment based on performance prediction: Towards an instance-aware problem solver. *Microsoft Research, Tech. Rep. MSR-TR-2005-125*, 2005.
- Frank Hutter, Youssef Hamadi, Holger H. Hoos, & Kevin Leyton-Brown. Performance prediction and automated tuning of randomized and parametric algorithms. In *International Conference on Principles and Practice of Constraint Programming*, pages 213–228. Springer, 2006.
- Frank Hutter, Domagoj Babic, Holger H. Hoos, & Alan J. Hu. Boosting verification by automatic tuning of decision procedures. In *Formal Methods in Computer Aided Design, FMCAD*, pages 27–34. IEEE, 2007a.
- Frank Hutter, Holger H. Hoos, & Thomas Stützle. Automatic algorithm configuration based on local search. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence, AAI*, pages 1152–1157. AAAI Press, 2007b.
- Frank Hutter, Holger H. Hoos, Kevin Leyton-Brown, & Kevin Murphy. An experimental investigation of model-based parameter optimisation: SPO and beyond. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 271–278. ACM, 2009a.
- Frank Hutter, Holger H. Hoos, Kevin Leyton-Brown, & Thomas Stützle. ParamILS: an automatic algorithm configuration framework. *Journal of Artificial Intelligence Research*, 36:267–306, 2009b.
- Frank Hutter, Holger H. Hoos, & Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration (extended version). *Technical Report TR-2010-10, University of British Columbia, Computer Science, Tech. Rep.*, 2010a.
- Frank Hutter, Holger H. Hoos, Kevin Leyton-Brown, & Kevin Murphy. Time-bounded sequential parameter optimization. In *International Conference on Learning and Intelligent Optimization*, pages 281–298. Springer, 2010b.
- Frank Hutter, Holger H. Hoos, & Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International conference on learning and intelligent optimization*, pages 507–523. Springer, 2011.
- Frank Hutter, Holger H. Hoos, & Kevin Leyton-Brown. Parallel algorithm configuration. In *International Conference on Learning and Intelligent Optimization*, pages 55–70. Springer, 2012.

- Frank Hutter, Holger H. Hoos, & Kevin Leyton-Brown. Bayesian optimization with censored response data. *arXiv preprint arXiv:1310.1947*, 2013.
- Frank Hutter, Manuel López-Ibáñez, Chris Fawcett, Marius Lindauer, Holger H. Hoos, Kevin Leyton-Brown, & Thomas Stützle. AClib: A benchmark library for algorithm configuration. In *International Conference on Learning and Intelligent Optimization*, pages 36–40. Springer, 2014a.
- Frank Hutter, Lin Xu, Holger H. Hoos, & Kevin Leyton-Brown. Algorithm runtime prediction: Methods & evaluation. *Artificial Intelligence*, 206:79–111, 2014b.
- Frank Hutter, Lars Kotthoff, & Joaquin Vanschoren, editors. *Automated Machine Learning - Methods, Systems, Challenges*. The Springer Series on Challenges in Machine Learning. Springer, 2019.
- Intel Corporation. Intel to Acquire SigOpt to Scale AI Productivity and Performance, Oct 2019. <https://www.intel.com/content/www/us/en/newsroom/news/sigopt-to-scale-ai-productivity-performance.html>. Accessed Nov. 29, 2021.
- Donald R. Jones. A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization*, 21(4):345–383, 2001.
- Donald R. Jones, Matthias Schonlau, & William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.
- Serdar Kadioglu, Yuri Malitsky, Meinolf Sellmann, & Kevin Tierney. ISAC - instance-specific algorithm configuration. In *19th European Conference on Artificial Intelligence, ECAI*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, pages 751–756. IOS Press, 2010.
- Giorgos Karafotias, Ágoston E. Eiben, & Mark Hoogendoorn. Generic parameter control with reinforcement learning. In *Genetic and Evolutionary Computation Conference, GECCO*, pages 1319–1326. ACM, 2014.
- Giorgos Karafotias, Mark Hoogendoorn, & Ágoston E. Eiben. Parameter control in evolutionary algorithms: Trends and challenges. *IEEE Transactions on Evolutionary Computation*, 19(2):167–187, 2015.
- Leonard Kaufman & Peter J. Rousseeuw. *Finding groups in data: an introduction to cluster analysis*, volume 344. John Wiley & Sons, 2009.
- Pascal Kerschke, Holger H. Hoos, Frank Neumann, & Heike Trautmann. Automated algorithm selection: Survey and perspectives. *Evolutionary computation*, 27(1):3–45, 2019.
- Jack Kiefer. Sequential minimax search for a maximum. *Proceedings of the American mathematical society*, 4(3):502–506, 1953.
- Robert Kleinberg, Kevin Leyton-Brown, & Brendan Lucier. Efficiency through procrastination: Approximately optimal algorithm configuration with runtime guarantees. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI*, pages 2023–2031. ijcai.org, 2017.

- Robert Kleinberg, Kevin Leyton-Brown, Brendan Lucier, & Devon Graham. Procrastinating with confidence: Near-optimal, anytime, adaptive algorithm configuration. *arXiv preprint arXiv:1902.05454*, 2019.
- Lars Kotthoff. Algorithm selection for combinatorial search problems: A survey. In *Data Mining and Constraint Programming*, pages 149–190. Springer, 2016.
- Christian Kroer & Yuri Malitsky. Feature filtering for instance-specific algorithm configuration. In *IEEE 23rd International Conference on Tools with Artificial Intelligence*, pages 849–855. IEEE, 2011.
- Michail G. Lagoudakis & Michael L. Littman. Algorithm selection using reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning, ICML*, pages 511–518. Morgan Kaufmann, 2000.
- Michail G. Lagoudakis & Michael L. Littman. Learning to select branching rules in the DPLL procedure for satisfiability. *Electronic Notes in Discrete Mathematics*, 9:344–359, 2001.
- Tor Lattimore & Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Hoong Chuin Lau & David Lo. Instance-based parameter tuning via search trajectory similarity clustering. In *International Conference on Learning and Intelligent Optimization*, pages 131–145. Springer, 2011.
- Kevin Leyton-Brown, Eugene Nudelman, & Yoav Shoham. Empirical hardness models: Methodology and a case study on combinatorial auctions. *Journal of the ACM, JACM*, 56(4):1–52, 2009.
- Marius Lindauer & Frank Hutter. Warmstarting of model-based algorithm configuration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018a.
- Marius Lindauer & Frank Hutter. Warmstarting of model-based algorithm configuration. In *Proceedings of the Thirty-Second Conference on Artificial Intelligence, AAAI*, pages 1355–1362, 2018b.
- Marius Lindauer, Frank Hutter, Holger H. Hoos, & Torsten Schaub. Autofolio: An automatically configured algorithm selector (extended abstract). In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI*, pages 5025–5029, 2017.
- Marius Lindauer, Katharina Eggensperger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, René Sass, & Frank Hutter. SMAC3: A versatile bayesian optimization package for hyperparameter optimization. *CoRR*, abs/2109.09831, 2021.
- Lindawati, Hoong Chuin Lau, & David Lo. Clustering of search trajectory and its application to parameter tuning. *Journal of the Operational Research Society*, 64(12):1742–1752, 2013a.

- Lindawati, Zhi Yuan, Hoong Chuin Lau, & Feida Zhu. Automated parameter tuning framework for heterogeneous and large instances: Case study in quadratic assignment problem. In *International Conference on Learning and Intelligent Optimization*, pages 423–437. Springer, 2013b.
- Lindawati, Feida Zhu, & Hoong Chuin Lau. FloTra: Flower-shape trajectory mining for instance-specific parameter tuning. In *Metaheuristics International Conference*. MIC, 2013c.
- Shengcai Liu, Ke Tang, & Xin Yao. Automatic construction of parallel portfolios via explicit instance grouping. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019*, pages 1560–1567. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33011560. URL <https://doi.org/10.1609/aaai.v33i01.33011560>.
- Shengcai Liu, Ke Tang, Yunwen Lei, & Xin Yao. On performance estimation in automatic algorithm configuration. In *The Thirty-Fourth Conference on Artificial Intelligence, AAAI*, pages 2384–2391. AAAI Press, 2020.
- Shengcai Liu, Ke Tang, & Xin Yao. Generative adversarial construction of parallel portfolios. *IEEE Trans. Cybern.*, 52(2):784–795, 2022. doi: 10.1109/TCYB.2020.2984546. URL <https://doi.org/10.1109/TCYB.2020.2984546>.
- Manuel Lopez-Ibanez & Thomas Stutzle. The automatic design of multiobjective ant colony optimization algorithms. *IEEE Transactions on Evolutionary Computation*, 16(6):861–875, 2012.
- Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Pérez Cáceres Leslie, Mauro Birattari, & Thomas Stützle. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3:43–58, 2016.
- Helena R. Lourenço, Olivier C. Martin, & Thomas Stützle. Iterated local search. In *Handbook of metaheuristics*, pages 320–353. Springer, 2003.
- R. Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.
- Gang Luo. A review of automatic selection methods for machine learning algorithms and hyper-parameter values. *Network Modeling Analysis in Health Informatics and Bioinformatics*, 5(1):1–16, 2016.
- Yuri Malitsky & Meinolf Sellmann. Instance-specific algorithm configuration as a method for non-model-based portfolio generation. In *International Conference on Integration of Artificial Intelligence and Operations Research*, pages 244–259. Springer, 2012.
- Yuri Malitsky, Deepak Mehta, & Barry O’Sullivan. Evolving instance specific algorithm configuration. In *Proceedings of the Sixth Annual Symposium on Combinatorial Search, SOCS*. AAAI Press, 2013a.

- Yuri Malitsky, Ashish Sabharwal, Horst Samulowitz, & Meinolf Sellmann. Algorithm portfolios based on cost-sensitive hierarchical clustering. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence, IJCAI*, pages 608–614. IJCAI, 2013b.
- Yuri Malitsky, Marius Merschformann, Barry O’Sullivan, & Kevin Tierney. Structure-preserving instance generation. In *International Conference on Learning and Intelligent Optimization*, pages 123–140. Springer, 2016.
- Oded Maron & Andrew W. Moore. Hoeffding races: Accelerating model selection search for classification and function approximation. In *Advances in Neural Information Processing Systems 6*, pages 59–66. Morgan Kaufmann, 1993.
- Joao Marques-Silva. Practical applications of boolean satisfiability. In *9th International Workshop on Discrete Event Systems*, pages 74–80. IEEE, 2008.
- Volodymyr Mnih, Csaba Szepesvári, & Jean-Yves Audibert. Empirical bernstein stopping. In *Machine Learning, Proceedings of the Twenty-Fifth International Conference, ICML*, volume 307 of *ACM International Conference Proceeding Series*, pages 672–679. ACM, 2008.
- Jonas Mockus. On bayesian methods for seeking the extremum. In *Optimization Techniques, IFIP Technical Conference*, pages 400–404, 1974.
- Andrew W. Moore & Mary S. Lee. Efficient algorithms for minimizing cross validation error. In *Machine Learning, Proceedings of the Eleventh International Conference*, pages 190–198. Morgan Kaufmann, 1994.
- Volker Nannen & A. E. Eiben. A method for parameter calibration and relevance estimation in evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO*, pages 183–190. ACM, 2006.
- Volker Nannen & A. E. Eiben. Efficient relevance estimation and value calibration of evolutionary algorithm parameters. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC*, pages 103–110. IEEE, 2007.
- Mihai Oltean. Evolving evolutionary algorithms using linear genetic programming. *Evolutionary Computation*, 13(3):387–410, 2005.
- OPTANO GmbH. OPTANO Algorithm Tuner v2.1.0, Nov 2021. <https://docs.optano.com/algorithm.tuner/2.1.0/>. Accessed Nov. 29, 2021.
- Ender Özcan, Mustafa Misir, Gabriela Ochoa, & Edmund K. Burke. A reinforcement learning - great-deluge hyper-heuristic for examination timetabling. *International Journal of Applied Metaheuristic Computing*, 1(1):39–59, 2010.
- Camille Pageau, Aymeric Blot, Holger H. Hoos, Marie-Eléonore Kessaci, & Laetitia Jourdan. Configuration of a dynamic mols algorithm for bi-objective flowshop scheduling. In *International Conference on Evolutionary Multi-Criterion Optimization*, pages 565–577. Springer, 2019.

- Leslie Pérez Cáceres, Manuel López-Ibáñez, Holger H. Hoos, & Thomas Stützle. An experimental study of adaptive capping in irace. In Roberto Battiti, Dmitri E. Kvasov, & Yaroslav D. Sergeyev, editors, *Learning and Intelligent Optimization, 11th International Conference, LION 11*, volume 10556 of *Lecture Notes in Computer Science*, pages 235–250. Springer, Cham, Switzerland, 2017. doi: 10.1007/978-3-319-69404-7\_17.
- James E. Pettinger & Richard M. Everson. Controlling genetic algorithms with reinforcement learning. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO*, page 692. Morgan Kaufmann, 2002.
- Robin L. Plackett. The analysis of permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 24(2):193–202, 1975.
- Yasha Pushak & Holger H. Hoos. Algorithm configuration landscapes. In *International Conference on Parallel Problem Solving from Nature*, pages 271–283. Springer, 2018.
- Yasha Pushak & Holger H. Hoos. Golden parameter search: exploiting structure to quickly configure parameters in parallel. In *Genetic and Evolutionary Computation Conference, GECCO*, pages 245–253. ACM, 2020.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley, 1994.
- William M Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336):846–850, 1971.
- Carl Edward Rasmussen & Christopher K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006.
- John R. Rice. The algorithm selection problem. *Advances in Computers*, 15:65–118, 1976.
- Enda Ridge & Daniel Kudenko. Tuning the performance of the MMAS heuristic. In *International Workshop on Engineering Stochastic Local Search Algorithms*, pages 46–60. Springer, 2007.
- Ranjit K. Roy. *A primer on the Taguchi method*. Society of Manufacturing Engineers, 2010.
- Rubén Ruiz & Concepción Maroto. A comprehensive review and evaluation of permutation flowshop heuristics. *European journal of operational research*, 165(2):479–494, 2005.
- S. J. Russell. Rationality and intelligence. *Artificial Intelligence*, 94:57–77, 1997.
- S.J. Russell. Rationality and intelligence: A brief update. In V. Müller, editor, *Fundamental Issues of Artificial Intelligence*. Springer, Cham., 2016.
- S.J. Russell & E.H. Wefald. *Do the Right Thing: Studies in Limited Rationality*. MIT Press, Cambridge, Massachusetts, 1991.
- Yoshitaka Sakurai, Kouhei Takada, Takashi Kawabe, & Setsuo Tsuruta. A method to control parameters of evolutionary algorithms by using reinforcement learning. In *Sixth International Conference on Signal-Image Technology and Internet-Based Systems, SITIS*, pages 74–79. IEEE Computer Society, 2010.

- Thomas J. Santner, Brian J. Williams, & William I. Notz. *The Design and Analysis of Computer Experiments*. Springer series in statistics. Springer, 2003. ISBN 978-0-387-95420-2.
- Robert E. Schapire. The boosting approach to machine learning: An overview. *Nonlinear estimation and classification*, pages 149–171, 2003.
- Meinolf Sellmann & Kevin Tierney. Hyper-parameterized dialectic search for non-linear box-constrained optimization with heterogenous variable types. In *Learning and Intelligent Optimization - 14th International Conference, LION*, volume 12096 of *Lecture Notes in Computer Science*, pages 102–116. Springer, 2020.
- Gresa Shala, André Biedenkapp, Noor Awad, Steven Adriaensen, Marius Lindauer, & Frank Hutter. Learning step-size adaptation in CMA-ES. In *International Conference on Parallel Problem Solving from Nature*, pages 691–706. Springer, 2020.
- Chiara F Sironi, Jialin Liu, & Mark HM Winands. Self-adaptive monte carlo tree search in general game playing. *IEEE Transactions on Games*, 12(2):132–144, 2018.
- Selmar K. Smit & A. E. Eiben. Comparing parameter tuning methods for evolutionary algorithms. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC*, pages 399–406. IEEE, 2009.
- Kate Smith-Miles & Simon Bowly. Generating new test instances by evolving in instance space. *Computers & Operations Research*, 63:102–113, 2015.
- David Speck, André Biedenkapp, Frank Hutter, Robert Mattmüller, & Marius Lindauer. Learning heuristic selection with dynamic algorithm configuration. In *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling, ICAPS*, pages 597–605. AAAI Press, 2021.
- Lee Spector & Eva Moscovici. Recent developments in autoconstructive evolution. In *Genetic and Evolutionary Computation Conference*, pages 1154–1156. ACM, 2017.
- Lee Spector & Alan J. Robinson. Genetic programming and autoconstructive evolution with the push programming language. *Genetic Programming and Evolvable Machines*, 3(1):7–40, 2002.
- Thomas Stützle & Manuel López-Ibáñez. Automated design of metaheuristic algorithms. In Michel Gendreau & Jean-Yves Potvin, editors, *Handbook of Metaheuristics*, volume 272 of *International Series in Operations Research & Management Science*, pages 541–579. Springer, 2019. doi: 10.1007/978-3-319-91086-4\_17.
- Ke Tang, Shengcai Liu, Peng Yang, & Xin Yao. Few-shots parallel algorithm portfolio construction via co-evolution. *IEEE Trans. Evol. Comput.*, 25(3):595–607, 2021. doi: 10.1109/TEVC.2021.3059661. URL <https://doi.org/10.1109/TEVC.2021.3059661>.
- Nguyen Dang Thi Thanh & Patrick De Causmaecker. Motivations for the development of a multi-objective algorithm configurator. In *Proceedings of the 3rd International Conference*

- on *Operations Research and Enterprise Systems, ICORES*, pages 328–333. SciTePress, 2014.
- Chris Thornton, Frank Hutter, Holger H. Hoos, & Kevin Leyton-Brown. Auto-weka: combined selection and hyperparameter optimization of classification algorithms. In Inderjit S. Dhillon, Yehuda Koren, Rayid Ghani, Ted E. Senator, Paul Bradley, Rakesh Parekh, Jingrui He, Robert L. Grossman, & Ramasamy Uthurusamy, editors, *The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013, Chicago, IL, USA, August 11-14, 2013*, pages 847–855. ACM, 2013. doi: 10.1145/2487575.2487629. URL <https://doi.org/10.1145/2487575.2487629>.
- Kevin Tierney & Yuri Malitsky. An algorithm selection benchmark of the container premarshalling problem. In *Learning and Intelligent Optimization - 9th International Conference, LION*, volume 8994 of *Lecture Notes in Computer Science*, pages 17–22. Springer, 2015.
- Alexander Tornede, Marcel Wever, & Eyke Hüllermeier. Towards meta-algorithm selection. *CoRR*, abs/2011.08784, 2020a. URL <https://arxiv.org/abs/2011.08784>.
- Alexander Tornede, Marcel Wever, & Eyke Hüllermeier. Extreme algorithm selection with dyadic feature representation. In *Discovery Science - 23rd International Conference, DS*, pages 309–324, 2020b.
- Alexander Tornede, Marcel Wever, Stefan Werner, Felix Mohr, & Eyke Hüllermeier. Run2survive: A decision-theoretic approach to algorithm selection based on survival analysis. In *Proceedings of The 12th Asian Conference on Machine Learning, ACML*, volume 129 of *Proceedings of Machine Learning Research*, pages 737–752. PMLR, 2020c.
- Alexander Tornede, Viktor Bengs, & Eyke Hüllermeier. Machine learning for online algorithm selection under censored feedback. *CoRR*, abs/2109.06234, 2021a.
- Alexander Tornede, Lukas Gehring, Tanja Tornede, Marcel Wever, & Eyke Hüllermeier. Algorithm selection on a meta level, 2021b.
- Leslie G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- Koen Van der Blom, Alex Serban, Holger H. Hoos, & Joost Visser. AutoML adoption in ml software. In *8th ICML Workshop on Automated Machine Learning (AutoML)*, 2021.
- Tim van Dijk, Martijn Mes, Marco Schutten, & Joaquim A. S. Gromicho. A unified race algorithm for offline parameter tuning. In *Proceedings of the 2014 Winter Simulation Conference*, pages 3971–3982. IEEE/ACM, 2014.
- Sander van Rijn, Carola Doerr, & Thomas Bäck. Towards an adaptive CMA-ES configurator. In *15th International Conference Parallel Problem Solving from Nature*, volume 11101 of *Lecture Notes in Computer Science*, pages 54–65. Springer, 2018.

- Diederick Vermetten, Sander van Rijn, Thomas Bäck, & Carola Doerr. Online selection of CMA-ES variants. In Anne Auger & Thomas Stützle, editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO*, pages 951–959. ACM, 2019.
- Ellen Vitercik. *Automated algorithm and mechanism configuration*. PhD thesis, Carnegie Mellon University, 2021.
- Dennis Wackerly, William Mendenhall, & Richard L. Scheaffer. *Mathematical statistics with applications*. Cengage Learning, 2014.
- Abraham Wald. *Sequential analysis*. Courier Corporation, 1947.
- Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, & Alexei A. Efros. Dataset distillation. *CoRR*, abs/1811.10959, 2018. URL <http://arxiv.org/abs/1811.10959>.
- Gellért Weisz, András György, & Csaba Szepesvári. LEAPSANDBOUNDS: A method for approximately optimal algorithm configuration. In *Proceedings of the 35th International Conference on Machine Learning, ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 5254–5262. PMLR, 2018.
- Gellért Weisz, András György, & Csaba Szepesvári. CapsAndRuns: An improved method for approximately optimal algorithm configuration. In *Proceedings of the 36th International Conference on Machine Learning, ICML*, volume 97, pages 6707–6715. PMLR, 2019.
- Gellért Weisz, András György, Wei-I Lin, Devon R. Graham, Kevin Leyton-Brown, Csaba Szepesvári, & Brendan Lucier. ImpatientCapsAndRuns: Approximately optimal algorithm configuration from an infinite pool. In *Annual Conference on Neural Information Processing Systems, NeurIPS*, 2020.
- Lin Xu, Frank Hutter, Holger H. Hoos, & Kevin Leyton-Brown. SATzilla: Portfolio-based algorithm selection for SAT. *Journal of artificial intelligence research*, 32:565–606, 2008.
- Lin Xu, Holger H. Hoos, & Kevin Leyton-Brown. Hydra: Automatically configuring algorithms for portfolio-based selection. In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence, AAAI*. AAAI Press, 2010.
- Lin Xu, Frank Hutter, Holger H. Hoos, & Kevin Leyton-Brown. Hydra-MIP: Automated algorithm configuration and selection for mixed integer programming. In *RCRA workshop on experimental evaluation of algorithms for solving problems with combinatorial explosion at the international joint conference on artificial intelligence, IJCAI*, pages 16–30, 2011.
- Mohamed El Yafrani, Marcella Scoczynski Ribeiro Martins, Inkyung Sung, Markus Wagner, Carola Doerr, & Peter Nielsen. MATE: A model-based algorithm tuning engine. *arXiv preprint arXiv:2004.12750*, 2020.
- Li Yang & Abdallah Shami. On hyperparameter optimization of machine learning algorithms: Theory and practice. *Neurocomputing*, 415:295–316, 2020.

- Furong Ye, Carola Doerr, & Thomas Bäck. Leveraging benchmarking data for informed one-shot dynamic algorithm selection. In *Genetic and Evolutionary Computation Conference, GECCO*, pages 245–246. ACM, 2021.
- Tong Yu & Hong Zhu. Hyper-parameter optimization: A review of algorithms and applications. *arXiv preprint arXiv:2003.05689*, 2020.
- Bo Yuan & Marcus Gallagher. A hybrid approach to parameter tuning in genetic algorithms. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC*, pages 1096–1103. IEEE, 2005.
- Bo Yuan & Marcus Gallagher. Combining Meta-EAs and racing for difficult EA parameter tuning tasks. In *Parameter setting in evolutionary algorithms*, pages 121–142. Springer, 2007.
- Zhi Yuan, Thomas Stützle, Marco Antonio Montes de Oca, Hoong Chuin Lau, & Mauro Birattari. An analysis of post-selection in automatic configuration. In *Genetic and Evolutionary Computation Conference, GECCO*, pages 1557–1564. ACM, 2013.
- Tiantian Zhang, Michael Georgiopoulos, & Georgios C. Anagnostopoulos. S-race: a multi-objective racing algorithm. In *Genetic and Evolutionary Computation Conference, GECCO*, pages 1565–1572. ACM, 2013.
- Tiantian Zhang, Michael Georgiopoulos, & Georgios C. Anagnostopoulos. Multi-objective model selection via racing. *IEEE transactions on cybernetics*, 46(8):1863–1876, 2015a.
- Tiantian Zhang, Michael Georgiopoulos, & Georgios C. Anagnostopoulos. SPRINT multi-objective model racing. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO*, pages 1383–1390. ACM, 2015b.
- Shlomo Zilberstein. Metareasoning and bounded rationality. In *Metareasoning - Thinking about Thinking*, pages 27–40. MIT Press, 2011.
- Marc-André Zöllner & Marco F. Huber. Benchmark and survey of automated machine learning frameworks. *J. Artif. Intell. Res.*, 70:409–472, 2021. doi: 10.1613/jair.1.11854. URL <https://doi.org/10.1613/jair.1.11854>.