

AC-Band: A Combinatorial Bandit-Based Approach to Algorithm Configuration

Jasmin Brandt¹, Elias Schede², Viktor Bengs³, Björn Haddenhorst¹,
Eyke Hüllermeier^{3,4}, Kevin Tierney²

¹Department of Computer Science, Paderborn University, Germany

²Decision and Operation Technologies Group, Bielefeld University, Germany

³Institute of Informatics, LMU Munich, Germany

⁴Munich Center for Machine Learning (MCML), Germany

Abstract

We study the algorithm configuration (AC) problem, in which one seeks to find an optimal parameter configuration of a given target algorithm in an automated way. Recently, there has been significant progress in designing AC approaches that satisfy strong theoretical guarantees. However, a significant gap still remains between the practical performance of these approaches and state-of-the-art heuristic methods. To this end, we introduce AC-Band, a general approach for the AC problem based on multi-armed bandits that provides theoretical guarantees while exhibiting strong practical performance. We show that AC-Band requires significantly less computation time than other AC approaches providing theoretical guarantees while still yielding high-quality configurations.

1 Introduction

Algorithm configuration (AC) is concerned with the task of automated search for a high-quality configuration (e.g., in terms of solution quality or runtime) of a given parameterized target algorithm. Parameterized algorithms are found in many different applications, including, for example, optimization (e.g., satisfiability (SAT) (Audemard and Simon 2018) or mixed-integer programming (MIP) solvers (IBM 2020)), simulation, and machine learning. Finding good configurations is a significant challenge for algorithm designers, as well as for users of such algorithms who may want to adjust the algorithm to perform well on data specific to their use case. Finding good configurations through trial and error by hand is a daunting task, hence automated AC methods have been developed on the basis of heuristics, such as ParamILS (Hutter, Hoos, and Stützle 2007; Hutter et al. 2009), GGA (Ansótegui, Sellmann, and Tierney 2009; Ansótegui et al. 2015), irace (Birattari et al. 2002; López-Ibáñez et al. 2016) or SMAC (Hutter, Hoos, and Leyton-Brown 2013, 2011).

While heuristic configurators have had great success at finding good configurations on a wide variety of target algorithms, they do not come with theoretical guarantees. To this end, the pioneering work of Kleinberg, Leyton-Brown, and Lucier (2017) proposes the first algorithm configurator with provable, theoretical guarantees on the (near-)optimality of the configuration returned. Further improvements and adjustments to those guarantees have followed (Weisz, György, and

Szepesvári 2018, 2019; Kleinberg et al. 2019; Weisz et al. 2020). All of these works essentially consider the runtime as the performance objective and provide, in particular, an upper bound on the total runtime of the respective algorithm configurator that is (nearly) optimal in a worst-case sense.

Despite their appealing theoretical properties and the steady progress on their empirical performance, these approaches still cannot compete with heuristic approaches in terms of practical performance. The main issue of these theoretical approaches is that they are conservative in the process of discarding configurations from the pool of potential candidates, as pointed out in recent work (Weisz et al. 2020). This is indeed a key characteristic difference compared to the heuristic approaches, which discard configurations quite quickly after being used only on a couple of problem instances. From a practical point of view, this makes sense, as the risk of discarding all good configurations is generally lower than wasting lots of time looking at bad configurations.

In an attempt to further bridge the gap between heuristic configurators and theoretical approaches, we propose AC-Band, a general algorithm configurator inspired by the popular Hyperband (Li et al. 2016) approach based on multi-armed bandits (Lattimore and Szepesvári 2020). Hyperband is an algorithm for the hyperparameter optimization problem (HPO) (Yang and Shami 2020; Bischl et al. 2021b), which is essentially a subproblem of the general AC problem focusing on configuring solution quality of algorithms, with a particular focus on machine learning methods. While using HPO approaches for AC looks attractive at first, it is rather uncommon in practice due to the subtle differences between the two problems. These differences include potentially using runtime as a configuration metric and the existence of multiple *problem instances*, which are different settings or scenarios that an optimization method must solve.

Our suggested approach reconciles the basic idea behind the mechanism of Hyperband with the key characteristics of the AC problem and incorporates the advantageous property of discarding configurations quickly. This is achieved by first replacing the underlying bandit algorithm of Hyperband, Successive Halving (SH) (Karnin, Koren, and Somekh 2013), by a more general variant, Combinatorial Successive Elimination (CSE) (Brandt et al. 2022), and then carefully adapting the parameters of the iterative CSE calls over time. In contrast to SH, as well as to all other multi-armed bandit algorithms

for pure exploration with finite budget, CSE allows (i) to steer the aggressiveness of discarding arms (configurations in our terminology) from the set of consideration, (ii) to pull more than one arm simultaneously (run multiple configurations in parallel), and (iii) to work with observations either of quantitative or qualitative nature. As mentioned above, the first property seems to be of major importance in AC problems, but the other two properties will turn out to be particularly helpful as well. Indeed, (ii) obviously allows parallelization of the search process, while the generality regarding the nature of the observations in (iii) transfers quite naturally to the suggested method.

The interplay of the second and third properties allows us to instantiate AC-Band to obtain appealing practical performance compared to existing theoretical approaches regarding the total runtime. On the theoretical side, we derive (under mild assumptions on the underlying AC problem) that AC-Band is guaranteed to return a nearly optimal configuration with high probability if used on sufficiently many problem instances. Our theoretical result is quite general in the sense that the notion of optimality is not restricted to the runtime of the configurations, but also encompasses other target metrics such as solution quality or memory usage.

2 Related Work

Theoretical advances in AC. The field of AC has grown to include many different methods and settings; we refer to Schede et al. (2022) for a full overview, especially with regard to the heuristic methods previously mentioned. Inspired by Kleinberg, Leyton-Brown, and Lucier (2017), who introduced *Structured Procrastination* together with a non-trivial worst-case runtime bound, more and more algorithms with even better theoretical guarantees with respect to the runtime have been proposed. *LeapsAndBounds* (Weisz, György, and Szepesvári 2018) tries to guess an upper bound on the optimal runtime by doubling the last failed guess, whereas *Structured Procrastination with confidence* (Kleinberg et al. 2019) works by delaying solving hard problem instances until later. Rather, it first runs the configurations with the smallest lower confidence bound of its mean runtime on instances that are easy to solve. Another recent method, *CapsAndRuns* (Weisz, György, and Szepesvári 2019), first estimates a timeout for each configuration and afterwards performs a Bernstein race over the configurations. As a follow up method, *Impatient-CapsAndRuns* (Weisz et al. 2020) uses a more aggressive elimination strategy by filtering unlikely optimal configurations in a preprocessing. Further theoretical progress has been made regarding the analysis of the estimation error in AC settings (Balcan et al. 2019; Balcan, Sandholm, and Vitercik 2020), the distribution of the computational budget (Liu et al. 2020) and the understanding of heuristic methods (Hall, Oliveto, and Sudholt 2019, 2020).

HPO. As a subset of AC, HPO involves setting the *hyperparameters* of algorithms, in particular machine learning approaches. The term hyperparameter differentiates parameters that change the behavior of the algorithm being configured from parameters that are induced or learned from data and are thus not set by a user. In contrast, AC focuses on configuring

algorithms that solve instances of a dataset independently. We refer to Bischl et al. (2021a) for a full overview of HPO.

Bandit methods for AC. Classically, methods for multi-armed bandits (MAB) (Lai and Robbins 1985; Bubeck and Cesa-Bianchi 2012; Lattimore and Szepesvári 2020) are designed to find a good balance between exploration-exploitation of specific choice alternatives (e.g., configurations or hyperparameters). The pure exploration setting, however, has attracted much research interest as well (Bubeck, Munos, and Stoltz 2009; Karnin, Koren, and Somekh 2013; Aziz et al. 2018), especially for HPO (Jamieson and Talwalkar 2015; Li et al. 2016). However, up to now bandit algorithms making single choice alternative decisions have been leveraged, although the parallel execution of configurations (or hyperparameters) in the AC (or HPO) setting seems to be predetermined for the combinatorial bandit variant (Cesa-Bianchi and Lugosi 2012; Chen, Wang, and Yuan 2013; Jourdan et al. 2021). In light of this, the recently proposed CSE algorithm (Brandt et al. 2022) seems promising as a generalization of the popular SH approach.

3 Problem Formulation

We adopt the formulation of the problem as by Schede et al. (2022). Let \mathcal{I} be the space of problem instances and \mathcal{P} an unknown probability distribution over \mathcal{I} . Suppose \mathcal{A} is an algorithm that can be run on any problem instance $i \in \mathcal{I}$, and has different parameters p_j coming from a known domain Θ_j for each $j \in \{1, \dots, m\}$. We call \mathcal{A} the *target algorithm* and the Cartesian product of its parameter domains $\Theta = \Theta_1 \times \dots \times \Theta_m$ the *configuration* or *search space* consisting of all feasible parameter *configurations*. For a configuration $\theta \in \Theta$, we denote by \mathcal{A}_θ an instantiation of the target algorithm \mathcal{A} with configuration θ . Running the target algorithm \mathcal{A} with configuration θ on a specific problem instance $i \in \mathcal{I}$ results in costs specified by an unknown, and possibly stochastic, function $c : \mathcal{I} \times \Theta \rightarrow \mathbb{R}$, i.e., $c(i, \theta)$ represents the costs of using \mathcal{A}_θ for problem instance i . Here, the costs can correspond to the runtime of \mathcal{A}_θ for i , but also to other relevant target metrics such as the solution quality or the memory usage.

Algorithm Configurator. The goal in algorithm configuration is, roughly speaking, to find a configuration that is optimal, or at least nearly optimal, with respect to the costs in a certain sense, which we specify below. The search for such configurations is achieved by designing an algorithm configurator \mathcal{AC} that (i) selects specific configurations in Θ and (ii) runs them on some (perhaps randomly) chosen problem instances in \mathcal{I} . To this end, the algorithm configurator uses a statistic $s : \bigcup_{t \in \mathbb{N}} \mathbb{R}^t \rightarrow \mathbb{R}$ that maps the observed cost of a configuration θ used for a set of problem instances i_1, \dots, i_t to a representative numerical value $s(c(i_1, \theta), \dots, c(i_t, \theta))$, that guides the search behavior of \mathcal{AC} . For example, s could be the arithmetic mean, i.e., $s(c(i_1, \theta), \dots, c(i_t, \theta)) = t^{-1} \sum_{s=1}^t c(i_s, \theta)$.

In this work, we are interested in algorithm configurators that can run several different configurations, up to a certain size k , in parallel on a selected problem instance. The algorithm configurator is given a fixed computational budget B ,

which represents the maximum number of such parallel runs and is set externally. For this purpose, let $\tilde{\Theta}_{[2,k]} = \{\tilde{\Theta} \subset \Theta \mid 2 \leq |\tilde{\Theta}| \leq k\}$ be the set of all possible subsets of parameter configurations that have at least size 2 and at most size k . Furthermore, let $\tilde{\Theta}_{[2,k]}(\theta) = \{\tilde{\Theta} \in \tilde{\Theta}_{[2,k]} \mid \theta \in \tilde{\Theta}\}$ be the set of all possible subsets of parameter configurations containing the configuration $\theta \in \Theta$. Note, that the observed cost $c(i, \theta)$ for running θ along with other configurations on an instance $i \in \mathcal{I}$ could depend on the respectively chosen configuration set $\tilde{\Theta} \in \tilde{\Theta}_{[2,k]}(\theta)$. For example, the algorithm configurator could stop the parallel solution process as soon as one of the configurations provides a solution, and set a default cost (penalty) for the configurations that did not complete. Hence, we write $c_{\tilde{\Theta}}$ for the cost function in the following to take this contingency into account. Finally, we introduce $s_{\theta|\tilde{\Theta}}(t) = s(c_{\tilde{\Theta}}(i_1, \theta), \dots, c_{\tilde{\Theta}}(i_t, \theta))$ which is the statistic of θ after running it in parallel with the configurations in $\tilde{\Theta} \setminus \{\theta\}$ on t problem instances i_1, \dots, i_t .

ϵ -optimal Configurations. Since the observed costs are potentially dependent on the chosen set of configurations to evaluate, we first introduce the following assumption on the limit behavior of the statistics.

$$(A1) : \forall \tilde{\Theta} \in \tilde{\Theta}_{[2,k]} \forall \theta \in \tilde{\Theta} : S_{\theta|\tilde{\Theta}} = \lim_{t \rightarrow \infty} s_{\theta|\tilde{\Theta}}(t) \text{ exists.}$$

In words, for each possible set of configurations the (possibly dependent) statistic of each configuration involved converges to some limit value if run on infinitely many problem instances. Recalling the example of s being the arithmetic mean, this is arguably a mild assumption and implicitly assumed by most approaches for AC problems, due to the considered i.i.d. setting. Since our assumption is more general, it would also allow considering non-stationary scenarios of AC.

The natural notion of an optimal configuration θ^* is a configuration that has the largest (configuration-set dependent) limit value over all configurations. Indeed, if we would replace $\tilde{\Theta}_{[2,k]}$ by the singleton sets of all configurations in Θ , this would correspond to the commonly used definition of the optimal configuration (see (Schede et al. 2022)), as the limit value would be then $\mathbb{E}[c(i, \theta)]^1$. However, in our case this notion of optimality has two decisive drawbacks, as first of all such a θ^* may not exist. Moreover, even if it exists, the search for it might be hopeless as the configuration space is infinite (or very large). The latter issue arises in the “usual” AC problem scenario as well, and is resolved by relaxing the objective to finding a “good enough” configuration. We adapt this notion of near optimality by resorting to the definition of an ϵ -best arm from the preference-based bandit literature (Bengs et al. 2021). For some fixed relaxation parameter $\epsilon > 0$, we call a configuration θ an ϵ -best configuration iff

$$\forall \tilde{\Theta} \in \tilde{\Theta}_{[2,k]}(\theta) : S_{\theta|\tilde{\Theta}} \geq S_{(1)|\tilde{\Theta}} - \epsilon, \quad (1)$$

where $S_{(1)|\tilde{\Theta}} \geq \dots \geq S_{(|\tilde{\Theta}|)|\tilde{\Theta}}$ is the ordering of $\{S_{\theta|\tilde{\Theta}}\}_{\theta \in \tilde{\Theta}}$. Although we have relaxed the notion of optimality, finding ϵ -best configurations is still often like searching for a needle in a haystack. Hence, we need to ensure that there is a

¹The expectation is w.r.t. \mathcal{P} and the possible randomness due to \mathcal{A}_θ and/or the cost generation.

sufficiently high probability that an ϵ -best configuration is included in a large random sample set of Θ :

$$(A2) : \text{the proportion of } \epsilon\text{-best configurations is } \alpha \in (0, 1).$$

Note this assumption is once again inspired by the bandit literature dealing with infinitely many arms (de Heide et al. 2021). By fixing the probability for the non-occurrence of an ϵ -best configuration to some $\delta \in (0, 1)$, Assumption (A2) ensures that a uniformly at random sampled set of configurations with size $N_{\alpha, \delta} = \lceil \log_{1-\alpha}(\delta) \rceil$ contains at least one ϵ -best configuration with probability at least $1 - \delta$.

Of course, an efficient algorithm configurator \mathcal{AC} that aims to find an ϵ -best configuration θ^* cannot verify the condition $S_{\theta^*|\tilde{\Theta}} \geq S_{(1)|\tilde{\Theta}} - \epsilon$ for every possible query set $\tilde{\Theta} \in \tilde{\Theta}_{[2,k]}(\theta^*)$, in particular when the number of configurations, and thus the cardinality of $\tilde{\Theta}_{[2,k]}(\theta^*)$, is infinite. Instead, \mathcal{AC} can only guarantee the above condition for a finite number of query sets and therefore it will always find a proxy for an ϵ -best configuration that does not have to be a true ϵ -best configuration. To guarantee that \mathcal{AC} finds a true ϵ -best configuration with high probability, we introduce the following assumption.

$$(A3) : \forall M \in \mathbb{N}, \forall \theta \in \Theta :$$

$$\begin{aligned} &\mathbb{P}\left(\forall i \in [M] : S_{\theta|\tilde{\Theta}_i} \geq S_{(1)|\tilde{\Theta}_i} - \epsilon\right) \\ &\Rightarrow (\forall \tilde{\Theta} \in \tilde{\Theta}_{[2,k]}(\theta) : S_{\theta|\tilde{\Theta}} \geq S_{(1)|\tilde{\Theta}} - \epsilon) \geq 1 - \psi(M), \end{aligned}$$

where $\tilde{\Theta}_1, \dots, \tilde{\Theta}_M \sim \text{Uniform}(\tilde{\Theta}_{[2,k]}(\theta))$ and $\psi : \mathbb{N} \rightarrow [0, 1]$ is a strictly monotone decreasing function. In words, the probability that a configuration θ is a global ϵ -best configuration increases with the number of (randomly chosen) configuration sets on which it is a local ϵ -best configuration, i.e., the characteristic condition in (1) is fulfilled.

Note that (A3) is a high-level assumption on the difficulty of the underlying AC problem. The “easier” the problem, the steeper the form of ψ and vice versa. As we do not impose further assumptions on ψ other than monotonicity, our theoretical results below are valid for a variety of AC problems.

4 AC-Band Algorithm

The AC-Band algorithm consists of iterative calls to CSE (Brandt et al. 2022) that allow it to successively reduce the size of a candidate set of configurations. We first describe how CSE works, then elaborate on its use in the AC-Band approach.

Combinatorial Successive Elimination. CSE (Algorithm 1) is a combinatorial bandit algorithm that, given a finite set of arms (configurations), finds an optimal arm² defined similarly to (1) for $\epsilon = 0$ using only a limited amount of feedback observations (budget). To this end, CSE proceeds on a round-by-round basis, in each of which (i) the non-eliminated arms are partitioned into groups of a given size k (line 3) if possible (lines 4–8, 18–19) and (ii) feedback for each group is queried under a round-specific budget, which, once exhausted, leads

²Its existence is simply assumed.

to the elimination of a certain fraction of arms in each group (Algorithm 2 called in lines 12 and 19). Here, the feedback observed for an arm is mapped to a numerical value using a statistic s that indicates the (group-dependent) utility of the arm, and is used as the basis for elimination in each round. The fraction of eliminated arms is steered via a function $f_\rho : [k] \rightarrow [k]$ with $f_\rho(x) = \lfloor x/2^\rho \rfloor$, where the predetermined parameter $\rho \in (0, \log_2(k))$ controls the aggressiveness of the elimination. A large value of ρ corresponds to a very aggressive elimination, retaining only the arm(s) with (the) highest statistics, while a small ρ eliminates only the arm(s) with the lowest statistics. The overall available budget is first split equally for each round, and then for all partitions in each round (line 2).

Algorithm 1: Combinatorial Successive Elimination (CSE)

Input: set of configurations $\tilde{\Theta}$ with $|\tilde{\Theta}| = n$, subset size $k \leq n$, budget B , $\rho \in (0, \log_2(k))$, problem instances I with $|I| = B$ which can be partitioned into $R^{\rho, k, n} = \min_x \{g^{\circ x}(n) \leq k\} + \min_x \{f_\rho^{\circ x}(k) \leq 1\}$ problem instances I_r with $|I_r| = P_r^{\rho, k, n} \cdot b_r$ where $\{P_r^{\rho, k, n}\}_r = \{\lfloor n/k (f_\rho(k)/k)^{r-1} \rfloor\}_r$, and $g(x) = f_\rho(k) \cdot \lfloor x/k \rfloor + x \bmod k$
Initialization: $\tilde{\Theta}_1 \leftarrow \tilde{\Theta}$, $r \leftarrow 1$

```

1: while  $|\tilde{\Theta}_r| \geq k$  do
2:    $b_r \leftarrow \lfloor B/(P_r^{\rho, k, n} \cdot R^{\rho, k, n}) \rfloor$ ,  $J \leftarrow P_r^{\rho, k, n}$ 
3:    $\tilde{\Theta}_{r,1}, \tilde{\Theta}_{r,2}, \dots, \tilde{\Theta}_{r,J} \leftarrow \text{Partition}(\tilde{\Theta}_r, k)$ 
4:   if  $|\tilde{\Theta}_{r,J}| < k$  then
5:      $\mathcal{R} \leftarrow \tilde{\Theta}_{r,J}$ ,  $J \leftarrow J - 1$ 
6:   else
7:      $\mathcal{R} \leftarrow \emptyset$ 
8:   end if
9:    $I_{r,1}, \dots, I_{r,J} \leftarrow \text{Partition}(I_r, b_r)$ 
10:   $\tilde{\Theta}_{r+1} \leftarrow \mathcal{R}$ 
11:  for  $j \in [J]$  do
12:     $\mathcal{R} \leftarrow \text{ArmElimination}(\tilde{\Theta}_{r,j}, b_r, f_\rho(|\tilde{\Theta}_{r,j}|), I_{r,j})$ 
13:     $\tilde{\Theta}_{r+1} \leftarrow \tilde{\Theta}_{r+1} \cup \mathcal{R}$ 
14:  end for
15:   $r \leftarrow r + 1$ 
16: end while
17:  $\tilde{\Theta}_{r+1} \leftarrow \emptyset$ 
18: while  $|\tilde{\Theta}_r| > 1$  do
19:    $\tilde{\Theta}_{r+1} \leftarrow \text{ArmElimination}(\tilde{\Theta}_r, b_r, f_\rho(|\tilde{\Theta}_r|), I_r)$ 
20:    $r \leftarrow r + 1$ 
21: end while

```

Output: The remaining item in $\tilde{\Theta}_r$

Algorithm 2: ArmElimination($\tilde{\Theta}'$, b , l , I')

```

1: Use  $\tilde{\Theta}'$  for  $b$  times on problem instances  $I'$ 
2: For all  $\theta \in \tilde{\Theta}'$ , update  $s_{\theta|\tilde{\Theta}'}(b)$ 
3: Choose an ordering  $\theta_1, \dots, \theta_{|\tilde{\Theta}'|}$  of  $(s_{\theta|\tilde{\Theta}'}(b))_{\theta \in \tilde{\Theta}'}$ 
4: Output:  $\{\theta_1, \dots, \theta_l\}$ 

```

In light of the AC problem we are facing, *querying feedback for a group of arms* corresponds to running a subset

of configurations in parallel on a problem instance, which results in observations in the form of costs. Moreover, we do not reuse a single problem instance for any parallel run so that the budget is in fact equal to the number of problem instances used. Accordingly, the overall budget of CSE corresponds to the number of problem instances used in total, which are split into disjoint problem instance sets of size b_r , i.e., the round-specific budget (line 9).

Since CSE initially assumes a finite set of arms and a fixed parameter ρ guiding the overall elimination aggressiveness, we face two trade-offs. The first is regarding the interplay between the number of initial arms and the round-wise budget:

(T1): If the initial number of configurations for CSE is small (large), the more (fewer) runs can be carried out on different instances, leading to potentially more reliable (unreliable) statistics, but only on a few (many) subsets of configurations.

The second trade-off arises through the interplay between the round-wise budget and the elimination aggressiveness:

(T2): If the elimination behavior of CSE is aggressive (conservative), then more (fewer) runs can be carried out on different instances, leading to potentially more reliable (unreliable) statistics, but only on a few (many) subsets of configurations.

The challenge now is to reconcile these two trade-offs and, above all, to take into account the specifics of AC problems.

Algorithm 3: AC-Band

Input: target algorithm \mathcal{A} , configuration space Θ , problem instance space \mathcal{I} , Budget B , subset size k , suboptimality ϵ of "good enough" configuration, proportion of ϵ -best configurations α , failure probability δ , $n_0 > \lceil \ln(\delta)/\ln(1-\alpha) \rceil \in \mathbb{N}$

Initialization: $E \leftarrow \lceil \log_2(n_0/n_0 - N_{\alpha, \delta}) \rceil$, $q \leftarrow 1 + k^{-1/E}$

$C_1 \leftarrow \log_q(2)$, $C_2 \leftarrow 1 + \log_q\left(n_0 + \frac{4n_0}{n_0 - N_{\alpha, \delta}}\right)$,

$C_3 \leftarrow \lceil \log_q(k) \rceil$

```

1: sample  $\theta_0 \in \Theta$ 
2: for  $e \in [E]$  do
3:    $n_e = \lceil n_0/2^e \rceil + 1$ ,  $\rho_e = \log_2(e+k-1/e)$ ,
4:   sample  $\theta_{e,1}, \dots, \theta_{e,n_e-1} \in \Theta$ 
5:   sample  $I_e \subset \mathcal{I} \setminus \cup_{e'=1}^{e-1} I_{e'}$  with  $|I_e| = B/c_e$ ,
     where  $c_e = \frac{(C_1 E - (2^{E-1} - 1)(2C_1 - C_2 - C_3))2^e}{2^E(-eC_1 + C_2 + C_3)}$ 
6:    $\theta_e = \text{CSE}(\{\theta_{e-1}, \theta_{e,1}, \dots, \theta_{e,n_e-1}\}, k, |I_e|, \rho_e, I_e)$ 
7: end for

```

Output: θ_E

AC-Band. The design of AC-Band (Algorithm 3) seeks to find a good balance for both tradeoffs (T1) and (T2) by calling CSE iteratively. Initially, CSE is invoked with larger sets of configurations, and an aggressive elimination strategy is applied. Over time, the size of the candidate sets is successively reduced, and the aggressiveness of the elimination strategy is also gradually decreased. Roughly speaking, the idea is to have high exploration in the beginning, and thus more risk that good configurations are discarded, and become more and more careful towards the end.

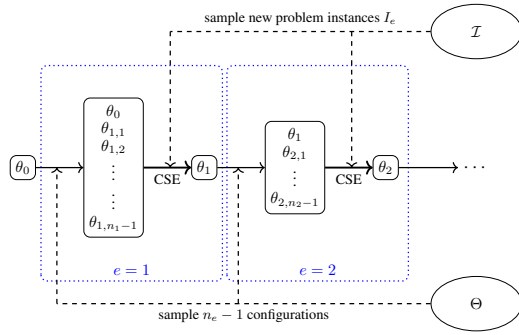


Figure 1: Illustration of AC-Band’s solving process.

More specifically, AC-Band proceeds in epochs $e \in \{1, \dots, E\}$, in each of which CSE is called on a specific set of problem instances using a specified degree of elimination aggressiveness ρ_e and a set of configurations of size n_e , with both ρ_e and n_e decreasing w.r.t. e (line 3). At the end of an epoch, i.e., when CSE terminates, a single high quality configuration among the considered set of configurations is returned (line 6). The set of configurations used in an epoch consists of the high quality configuration of the previous epoch, which is sampled randomly for the first epoch (line 1), and $n_e - 1$ randomly sampled configurations, which have not been considered before (line 4).

This epoch-wise procedure of AC-Band is depicted in Figure 1. Although AC-Band is similar in design to Hyperband in that it tries to find a good balance between specific trade-offs by successively invoking a bandit algorithm (CSE vs. SH), AC-Band differs in the way it defines the quality of the search objects³. Unlike Hyperband, we do not run configurations individually on problem instances and consider the cost of each configuration on its own, rather configurations are run in parallel and we consider potential interactions. Accordingly, we do not have one global quality value that we can compare for all configurations seen, but several at once.

The overall number of considered problem instances is B (the evaluation budget), which is a parameter that we analyze below. Besides the “usual” parameters of an algorithm configurator, i.e., the target algorithm \mathcal{A} , its configuration space Θ and the problem instance space \mathcal{I} , AC-Band requires:

- the maximum number of configurations that can be run in parallel k ,
- some relevant summary statistic $s : \bigcup_{t \in \mathbb{N}} \mathbb{R}^t \rightarrow \mathbb{R}$ for the observed costs (see Section 3),
- the theoretically motivated guarantee parameters $\epsilon > 0$, $\alpha \in (0, 1)$, and $\delta \in (0, 1)$ (see Section 3)
- a reference size n_0 for the set of epoch-wise sampled configurations (must be $\geq \lceil \frac{\ln(\delta)}{\ln(1-\alpha)} \rceil$ for technical reasons).

With these parameters specified, AC-Band determines the overall number of epochs E and the sufficient number of problem instances B/c_e (line 5) for an epoch-wise CSE to return a high quality configuration (see Section 5). Moreover, the overall number of considered configurations is guaranteed

³Hyperband uses hyperparameters of machine learning models as search objects and AC-Band algorithm configurations.

to be at least $N_{\alpha, \delta} = \lceil \frac{\ln(\delta)}{\ln(1-\alpha)} \rceil$, which in light of the random sampling of the configurations ensures that at least one ϵ -best configuration is sampled with a probability of at least $1 - \delta$.

5 Theoretical Guarantees

In Appendix C, we prove the following theoretical guarantee for AC-Band regarding the sufficient evaluation budget (or number of problem instances) to find an ϵ -best configuration with high probability w.r.t. \mathcal{P} as well as the randomness invoked by AC-Band. For the proof, we need to extend the theoretical guarantees for CSE to the setting of finding ϵ -best configurations (see Appendix B).

Theorem 5.1. *Let R^e be the number of rounds of CSE in epoch $e \in \{1, \dots, E\}$, let C_1, C_2 and C_3 be as defined in Alg. 3 and further let $\mathbb{A}_r(\theta^*)^4$ be the partition in round r of CSE containing θ^* and*

$$\bar{\gamma}^{-1} = \max_{e \in [E], r \in [R^e]} \left(1 + \bar{\gamma}_{\mathbb{A}_r(\theta^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^e]} \frac{\Delta(f_{\rho}(|\mathbb{A}_r(\theta^*)|)+1)|_{\mathbb{A}_r(\theta^*)}}{2} \right\} \right) \right),$$

$$\Delta_{(i)|\bar{\Theta}} = S_{(i)|\bar{\Theta}} - S_{(i)|\bar{\Theta}},$$

$$\bar{\gamma}_{\mathbb{A}_r(\theta^*)}^{-1}(t) = \min_{\theta \in \mathbb{A}_r(\theta^*)} \inf_{t' \in \mathbb{N}} \{ |s_{\theta|\mathbb{A}_r(\theta^*)}(t') - S_{\theta|\mathbb{A}_r(\theta^*)}| \leq t \}.$$

Under Assumptions (A1)–(A3), Algorithm 3 used with a subset size k , an ϵ -best configuration proportion of α , and a failure probability of δ finds an ϵ -best configuration θ^ with probability at least $\min\{1 - \delta, 1 - \psi((R^E)^{-1})\}$ if*

$$B \geq \bar{\gamma}^{-1} \cdot \frac{n_0}{k} \cdot \frac{C_1 E - (2^E - 1)(2C_1 - C_2 - C_3)}{2^E}.$$

Roughly speaking, AC-Band finds a near-optimal configuration with a probability depending on the allowed failure probability δ and the probability that a locally optimal configuration is also a globally optimal one (see Assumption (A3)) if the budget is large enough. The sufficient budget, in turn, is essentially driven by $\bar{\gamma}^{-1}$, which depends on the difficulty of the underlying AC problem by two characteristics: the maximal inverse convergence speed $\bar{\gamma}_{\mathbb{A}_r(\theta^*)}^{-1}$ of the used statistic s , and the maximal (halved) suboptimality gap Δ of the limits of the statistic between the best configuration and the best one that will be discarded from the query set θ^* is contained. The remaining terms of the sufficient budget can be computed explicitly once the theoretical guarantee parameters α and δ , as well as the subset size k , are fixed. The sufficient budget in dependence of the mentioned parameters is discussed and plotted in Appendix C.3.

Note that the theoretical guarantee in Theorem 5.1 is not directly comparable to the ones by the theoretical AC approaches (Kleinberg et al. 2019; Weisz, György, and Szepesvári 2018, 2019; Weisz et al. 2020). This is due to the major differences of our approach and the later ones on how we approach the AC problem. Indeed, we do not restrict ourselves to runtime as the target metric (or the costs), and

⁴ $\mathbb{A}_r(\theta) = \emptyset$ if θ is not contained anymore in the set of active configurations in round r .

we also take possible dependencies in the parallel runs into account. As a consequence, the notion of near optimality of a configuration in the other works is tailored more towards runtimes in an absolute sense, i.e., without considering interaction effects, while ours is more general and in particular focusing on such interaction effects. Thus, the theoretical guarantee in Theorem 5.1 in the form of a sufficient evaluation budget to obtain a nearly optimal configuration is sensible, as we do not commit to a specific target metric.

6 Experiments

We examine AC-Band on several standard datasets for evaluating theoretical approaches for AC. We note that while these datasets refer exclusively to runtimes, AC-Band is applicable to other target metrics (see Section 3). In our experiments, we address the following two research questions: **Q1**: Is AC-Band able to find high quality configurations in less CPU time than state-of-the-art AC methods with guarantees? **Q2**: How does AC-Band scale with k ?

Datasets. We use one SAT and two MIP datasets that have been used before to assess theoretical works on AC (Weisz et al. 2020). Due to space constraints, we only consider one of the MIP datasets here, while the appendix also discusses the other. The SAT dataset contains precomputed runtimes of configurations of the MiniSat SAT solver (Eén and Sörensson 2003) obtained by solving instances that are generated using CNFuzzDD (Weisz, György, and Szepesvári 2018). The dataset contains runtimes for 948 configurations on 20118 instances. The MIP datasets curated by Weisz et al. (2020) are generated using an Empirical Performance Model (EPM)(Hutter et al. 2014). In particular, the EPM is trained on the CPLEX solver (IBM 2020) separately using instances from a combinatorial auction (Regions200 (Leyton-Brown, Pearson, and Shoham 2000)) and wildlife conservation (RCW (Ahmadizadeh et al. 2010)) dataset. The resulting model is used to predict runtimes for 2000 configurations and 50000 and 35640 new instances, respectively. Since all runtimes are precomputed (a timeout of 900 seconds is used), the evaluation of configuration-instance pairs only required table look-ups in our experiments.

Evaluation. To compare methods, we consider two metrics: (i) the accumulated CPU time needed by a method to find a configuration, and (ii) the percent gap to the best configuration. This second metric measures in percent how much more time the configuration returned by the method needs to solve all instances compared to the best overall configuration for the dataset. Smaller values indicate that the configuration found is closer to the best configuration in terms of runtime and the best configuration has a value of 0. This allows for comparing the quality between methods, as well as to determine how “far” a configuration is from the optimal one. In practical applications of AC, wall-clock time is often a bottleneck, and speeding up the process of finding a suitable configuration is the main focus. For these speedups, practitioners are (usually) willing to sacrifice configuration quality to a certain extent. The other theoretical works use the R^δ metric (note that δ has a different meaning in this work) to evaluate the quality of a returned configuration. This metric is a variation of the mean runtime, where the mean runtime

of a configuration is only computed over the $(1 - \delta)$ portion of instances with the lowest runtime. In real-world settings, we do not have the luxury of ignoring a part of the instance set, thus we do not view this metric as suitable for evaluating our approach. For the sake of completeness, we nevertheless report the R^δ values in Appendix D.

Initialization of AC-Band. Due to the generality of our approach, a summary statistic s that measures the quality of a configuration needs to be determined. In our case, the k configurations in a subset of CSE can be evaluated in parallel for an instance given that k CPUs are available. When running k configurations in parallel, time can be saved by stopping all remaining configuration runs as soon as the first configuration finishes on the given instance. Through this capping, we obtain right-censored feedback where a runtime is only observed for the “finisher”. A statistic that is able to deal with this censored feedback is needed to avoid using an imputation technique that could potentially add bias to the feedback. In line with Brandt et al. (2022), we interpret the obtained feedback as a preference: the finishing configuration is preferred over the remaining, unfinished configurations. Once we have obtained these preferences for multiple instances, we can use a preference-based statistic such as the relative frequency to establish an ordering over the configurations in k . In particular, we count how many times a configuration finishes first over a set of instances⁵.

Competitors. AC-Band is compared against ICAR, CAR++ (Weisz et al. 2020) and Hyperband (Li et al. 2016). At the moment, ICAR is the best performing AC method that comes with theoretical guarantees. We use the implementation provided by Weisz et al. (2020) with the same hyperparameter settings. Since AC-Band is inspired by Hyperband, we also adapt Hyperband for the AC setting for a comparison. Specifically, we set the parameter R of Hyperband such that it uses the same total budget (number of instances) as AC-Band. In addition, we use the average runtime over instances as the validation loss and consequently return the configuration with smallest average runtime. Finally, we set $s_{max} = \lceil (\log_\eta(n_{max})) \rceil$, adjust the calculation of r_i slightly to account for instances that were already seen, and try different values of η .

Choice of δ . Varying δ can lead to significantly different performance of AC-Band and other techniques. Due to space constraints, we only show results for two datasets for $\delta = 0.05$ since this setting has also been used in previous work (Weisz et al. 2020). We note, however, that other settings are just as valid, and therefore also provide results for $\delta = 0.01$ and additional datasets in Appendix D. In fact, using $\delta = 0.01$ can result in finding better configurations, albeit it is up to the user of the AC approach to decide which δ best suits their needs.

Q1 Figure 2 shows the CPU time used by each method and the percent gap to best configuration. With a small value of k , AC-Band lies on the Pareto front of percent gap versus CPU time for both datasets (for the third, Regions200, as well). This allows us to answer **Q1** in the affirmative. In particular, with $k = 2$ AC-Band is 72% percent faster than ICAR and

⁵Code: <https://github.com/DOTBielefeld/ACBand>




Figure 2: Mean CPU time and percent gap to best over 5 seeds for $\delta = 0.05$ and different α (columns) for AC-Band, ICAR, CAR++ and Hyperband on CNFuzzDD (top) and RCW (bottom). Circles indicate variants of AC-Band. Rectangles represent the standard error over the seeds. The number of configurations tried for CAR++: {97, 245, 492}, ICAR: {134, 351, 724}, AC-Band: {60, 153, 303}, Hyperband($\eta = 5$): {842}, Hyperband($\eta = 8$): {618}.

73% faster than Hyperband for $\delta = 0.05$ over all α and datasets, while providing configurations that are only 7% and 6% worse in terms of the gap to the best configuration. For most real-world applications, this is an acceptable trade-off in time versus quality. For all datasets, the percent gap to best decreases when AC-Band samples more configurations (smaller α). This increase in exploration does not lead to a significant increase in CPU time for a fixed k , since AC-Band still has the same budget, i.e., additional configurations are evaluated on fewer instances.

Hyperband samples more configurations in total than AC-Band, which helps it to achieve a better percent gap to best on datasets where the majority of configurations have a high mean runtime. On these datasets, only a few good configurations are present. This is the case for the RCW dataset (and the Regions200 dataset in the appendix) where only a few instances are needed to determine that a configuration should be discarded. On datasets where the runtime of configurations is more evenly distributed, such as the CNFuzzDD dataset, using too few instances may lead to discarding promising configurations early, giving AC-Band an edge by evaluating less configurations more thorough. Lastly, since Hyperband does not use capping, its CPU time deteriorates.

Q2 Our experiments clearly indicate that lower values of k are preferable. With $k = 2$, more CSE rounds are performed and thus the number of configurations decreases slower than with a higher k . With a larger k , the opposite occurs, and significant amounts of CPU time are expended with little information gain. However, note that higher k 's have a lower wall-clock time, so a user would receive answers sooner.

7 Conclusion

In this paper we introduced AC-Band, a versatile approach for AC problems that comes with theoretical guarantees even for a range of target metrics in addition to runtime. We showed that AC-Band returns a nearly optimal configuration w.r.t. the target metric with high probability if used with a sufficient number of problem instances and the underlying AC problem satisfies some mild assumptions. In our experimental study, we considered an instantiation of AC-Band based on preference feedback, which generally leads to faster average CPU times than other theoretical approaches, while still returning a suitable configuration in the end.

Our results open up several possibilities for future work. First, it would be interesting to analyze AC-Band specifically for the case in which runtime is the relevant target metric and investigate whether a similar worst-case overall runtime guarantee can be derived as for the other theoretical approaches in this vein. Next, a theoretical as well as empirical analysis regarding the interplay between the explicit instantiation of AC-Band w.r.t. the underlying statistic s and the characteristics of the underlying AC would be desirable. In other words, on what types of AC problems does a specific instantiation of AC perform well or poorly? Also, our mild Assumption (A1) would even allow some leeway in the configuration or problem instance sampling strategy of the algorithm configurator, which is currently simply uniformly at random for AC-Band. Finally, real-world AC applications generally have only a handful of instances available, thus it would be advantageous to have strong theoretical guarantees even for scenarios without thousands of instances.

Acknowledgements

This research was supported by the research training group Dataninja (Trustworthy AI for Seamless Problem Solving: Next Generation Intelligence Joins Robust Data Analysis) funded by the German federal state of North Rhine-Westphalia and by the German Research Foundation (DFG) within the project “Online Preference Learning with Bandit Algorithms” (project no. 317046553).

References

- Ahmadizadeh, K.; Dilkina, B.; Gomes, C. P.; and Sabharwal, A. 2010. An empirical study of optimization for maximizing diffusion in networks. In *International Conference on Principles and Practice of Constraint Programming*, 514–521. Springer.
- Ansótegui, C.; Malitsky, Y.; Samulowitz, H.; Sellmann, M.; and Tierney, K. 2015. Model-Based Genetic Algorithms for Algorithm Configuration. In *IJCAI*.
- Ansótegui, C.; Sellmann, M.; and Tierney, K. 2009. A Gender-Based Genetic Algorithm for the Automatic Configuration of Algorithms. 142–157.
- Audemard, G.; and Simon, L. 2018. On the glucose SAT solver. *International Journal on Artificial Intelligence Tools*, 27(01): 1840001.
- Aziz, M.; Anderton, J.; Kaufmann, E.; and Aslam, J. 2018. Pure exploration in infinitely-armed bandit models with fixed-confidence. In *Algorithmic Learning Theory*, 3–24. PMLR.
- Balcan, M.; DeBlasio, D. F.; Dick, T.; Kingsford, C.; Sandholm, T.; and Vitercik, E. 2019. How much data is sufficient to learn high-performing algorithms? *CoRR*, abs/1908.02894.
- Balcan, M.; Sandholm, T.; and Vitercik, E. 2020. 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, 580–590. PMLR.
- Bengs, V.; Busa-Fekete, R.; El Mesaoudi-Paul, A.; and Hüllermeier, E. 2021. Preference-based Online Learning with Dueling Bandits: A Survey. *Journal of Machine Learning Research*, 22: 1–108.
- Birattari, M.; Stützle, T.; Paquete, L.; and Varrentrapp, K. 2002. A Racing Algorithm for Configuring Metaheuristics. 11–18.
- Bischl, B.; Binder, M.; Lang, M.; Pielok, T.; Richter, J.; Coors, S.; Thomas, J.; Ullmann, T.; Becker, M.; Boulesteix, A.-L.; Deng, D.; and Lindauer, M. 2021a. Hyperparameter Optimization: Foundations, Algorithms, Best Practices and Open Challenges.
- Bischl, B.; Binder, M.; Lang, M.; Pielok, T.; Richter, J.; Coors, S.; Thomas, J.; Ullmann, T.; Becker, M.; Boulesteix, A.-L.; et al. 2021b. Hyperparameter optimization: Foundations, algorithms, best practices and open challenges. *arXiv preprint arXiv:2107.05847*.
- Brandt, J.; Haddenhorst, B.; Bengs, V.; and Hüllermeier, E. 2022. Finding Optimal Arms in Non-stochastic Combinatorial Bandits with Semi-bandit Feedback and Finite Budget.
- Bubeck, S.; and Cesa-Bianchi, N. 2012. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *CoRR*, abs/1204.5721.
- Bubeck, S.; Munos, R.; and Stoltz, G. 2009. Pure exploration in multi-armed bandits problems. In *International conference on Algorithmic learning theory*, 23–37. Springer.
- Cesa-Bianchi, N.; and Lugosi, G. 2012. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5): 1404–1422.
- Chen, W.; Wang, Y.; and Yuan, Y. 2013. Combinatorial Multi-Armed Bandit: General Framework and Applications. In *Proceedings of the 30th International Conference on Machine Learning*, 151–159.
- de Heide, R.; Cheshire, J.; Ménard, P.; and Carpentier, A. 2021. Bandits with many optimal arms. In *Advances in Neural Information Processing Systems*, volume 34, 22457–22469. Curran Associates, Inc.
- Eén, N.; and Sörensson, N. 2003. An extensible SAT-solver. In *International conference on theory and applications of satisfiability testing*, 502–518. Springer.
- Hall, G. T.; Oliveto, P. S.; and Sudholt, D. 2019. On the Impact of the Cutoff Time on the Performance of Algorithm Configurators. *CoRR*, abs/1904.06230.
- Hall, G. T.; Oliveto, P. S.; and Sudholt, D. 2020. Analysis of the Performance of Algorithm Configurators for Search Heuristics with Global Mutation Operators. *CoRR*, abs/2004.04519.
- Hutter, F.; Hoos, H.; and Leyton-Brown, K. 2013. Bayesian Optimization With Censored Response Data.
- Hutter, F.; Hoos, H.; Leyton-Brown, K.; and Stützle, T. 2009. ParamILS: An Automatic Algorithm Configuration Framework. *J. Artif. Intell. Res. (JAIR)*, 36: 267–306.
- Hutter, F.; Hoos, H.; and Stützle, T. 2007. Automatic Algorithm Configuration based on Local Search.
- Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2011. Sequential Model-Based Optimization for General Algorithm Configuration. In *Learning and Intelligent Optimization*, 507–523. Springer Berlin Heidelberg.
- Hutter, F.; Xu, L.; Hoos, H. H.; and Leyton-Brown, K. 2014. Algorithm runtime prediction: Methods & evaluation. *Artificial Intelligence*, 206: 79–111.
- IBM. 2020. ILOG CPLEX Optimization Studio 20.1.0: User’s Manual.
- Jamieson, K. G.; and Talwalkar, A. 2015. Non-stochastic Best Arm Identification and Hyperparameter Optimization. *CoRR*, abs/1502.07943.
- Jourdan, M.; Mutný, M.; Kirschner, J.; and Krause, A. 2021. Efficient Pure Exploration for Combinatorial Bandits with Semi-Bandit Feedback.
- Karnin, Z.; Koren, T.; and Somekh, O. 2013. Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, 1238–1246. PMLR.
- Kleinberg, R.; Leyton-Brown, K.; and Lucier, B. 2017. Efficiency Through Procrastination: Approximately Optimal

- Algorithm Configuration with Runtime Guarantees. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI, 2023–2031*. ijcai.org.
- Kleinberg, R.; Leyton-Brown, K.; Lucier, B.; and Graham, D. 2019. Procrastinating with confidence: Near-optimal, anytime, adaptive algorithm configuration. *arXiv preprint arXiv:1902.05454*.
- Lai, T.; and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1): 4–22.
- Lattimore, T.; and Szepesvári, C. 2020. *Bandit algorithms*. Cambridge University Press.
- Leyton-Brown, K.; Pearson, M.; and Shoham, Y. 2000. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings of the 2nd ACM conference on Electronic commerce*, 66–76.
- Li, L.; Jamieson, K. G.; DeSalvo, G.; Rostamizadeh, A.; and Talwalkar, A. 2016. Efficient Hyperparameter Optimization and Infinitely Many Armed Bandits. *CoRR*, abs/1603.06560.
- Liu, S.; Tang, K.; Lei, Y.; and Yao, X. 2020. On Performance Estimation in Automatic Algorithm Configuration. In *The Thirty-Fourth Conference on Artificial Intelligence, AAAI, 2384–2391*. AAAI Press.
- López-Ibáñez, M.; Dubois-Lacoste, J.; Pérez Cáceres, L.; Birattari, M.; and Stützle, T. 2016. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3: 43–58.
- Schede, E.; Brandt, J.; Tornede, A.; Wever, M.; Bengs, V.; Hüllermeier, E.; and Tierney, K. 2022. A Survey of Methods for Automated Algorithm Configuration. *CoRR*, abs/2202.01651.
- Weisz, G.; György, A.; Lin, W.; Graham, D. R.; Leyton-Brown, K.; Szepesvári, C.; and Lucier, B. 2020. Impatient-CapsAndRuns: Approximately Optimal Algorithm Configuration from an Infinite Pool. In *Annual Conference on Neural Information Processing Systems, NeurIPS*.
- Weisz, G.; György, A.; and Szepesvári, C. 2018. LEAP-SANDBOUNDS: 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, 5254–5262*. PMLR.
- Weisz, G.; György, A.; and Szepesvári, C. 2019. CapsAndRuns: An Improved Method for Approximately Optimal Algorithm Configuration. In *Proceedings of the 36th International Conference on Machine Learning, ICML, volume 97, 6707–6715*. PMLR.
- Yang, L.; and Shami, A. 2020. On hyperparameter optimization of machine learning algorithms: Theory and practice. *Neurocomputing*, 415: 295–316.

A List of Symbols

The following table contains a list of symbols that are frequently used in the main paper as well as in the following supplementary material.

Basics	
$\mathbb{1}\{\cdot\}$	Indicator function
\mathbb{N}	Set of natural numbers (without 0), i.e., $\mathbb{N} = \{1, 2, 3, \dots\}$
\mathbb{R}	Set of real numbers
\mathcal{I}	Space of problem instances
\mathcal{P}	Probability distribution over instance space \mathcal{I}
\mathcal{A}	Target algorithm that can be run on problem instance $i \in \mathcal{I}$
Θ	Configuration search space consisting all feasible parameter configurations
\mathcal{A}_θ	Instantiation of \mathcal{A} with configuration θ
$c(i, \theta)$	$c : \mathcal{I} \times \Theta \rightarrow \mathbb{R}$ costs of using \mathcal{A}_θ on problem instance $i \in \mathcal{I}$
\mathcal{AC}	Algorithm configurator
k	Maximal possible subset size
B	Budget for the learner (algorithm configurator)
$\Theta_{[2,k]}$	All subsets of Θ of size at least 2 and at most k : $\{\tilde{\Theta} \subseteq \Theta \mid 2 \leq \tilde{\Theta} \leq k\}$
$\Theta_{[2,k]}(\theta)$	All subsets in $\Theta_{[2,k]}$ which contain configuration θ : $\{\tilde{\Theta} \in \Theta_{[2,k]} \mid \theta \in \tilde{\Theta}\}$
$c_{\tilde{\Theta}}$	Cost of running θ on problem instance $i \in \mathcal{I}$ in parallel with configurations $\theta' \in \tilde{\Theta} \in \Theta_{[2,k]} \setminus \{\theta\}$
Modelling related	
s	Statistic mapping costs to a numerical value: $s : \cup_{t \in \mathbb{N}} \mathbb{R}^t \rightarrow \mathbb{R}, (c(i_1, \theta), \dots, c(i_t, \theta)) \mapsto s(c(i_1, \theta), \dots, c(i_t, \theta))$
$s_{\theta \tilde{\Theta}}(t)$	$s_{\theta \tilde{\Theta}}(t) = s(c_{\tilde{\Theta}}(i_1, \theta), \dots, c_{\tilde{\Theta}}(i_t, \theta))$ statistic of $\theta \in \Theta$ after running $\tilde{\Theta}$ in parallel
$S_{\theta \tilde{\Theta}}$	Limit of the statistics for configuration θ in query set $\tilde{\Theta}$, i.e., $\lim_{t \rightarrow \infty} s_{\theta \tilde{\Theta}}(t)$ (see Assumption (A1))
ϵ	Near-optimality parameter, $\epsilon \in (0, 1)$
α	Proportion of ϵ -best configurations in Θ , $\alpha \in (0, 1)$ (see Assumption (A2))
δ	Fixed error probability for identifying an ϵ -best configuration, $\delta \in (0, 1)$
$N_{\alpha, \delta}$	$N_{\alpha, \delta} = \lceil \log_{1-\alpha}(\delta) \rceil$ number of configurations that have to be sampled to ensure that at least one ϵ -best configuration is contained with probability at least $1 - \delta$
$\gamma_{\theta \tilde{\Theta}}$	Pointwise minimal function such that $ s_{\theta \tilde{\Theta}}(t) - S_{\theta \tilde{\Theta}} \leq \gamma_{\theta \tilde{\Theta}}(t)$ for all t
$\gamma_{\theta \tilde{\Theta}}^{-1}(t)$	$\min\{t' \in \mathbb{N} \mid s_{\theta \tilde{\Theta}}(t') - S_{\theta \tilde{\Theta}} \leq t\}$ (exists due to Assumption (A1))
$\bar{\gamma}_{\tilde{\Theta}}^{-1}(t)$	Minimal $\gamma_{\theta \tilde{\Theta}}^{-1}(t) = \min\{t' \in \mathbb{N} \mid s_{\theta \tilde{\Theta}}(t') - S_{\theta \tilde{\Theta}} \leq t\}$ over all $\theta \in \tilde{\Theta}$ (exists due to Assumption (A1))
$S_{(l) \tilde{\Theta}}, \Delta_{(l) \tilde{\Theta}}$	l -th order statistic of $\{S_{\theta \tilde{\Theta}}\}_{\theta \in \tilde{\Theta}}$ for $l \in \{1, 2, \dots, \tilde{\Theta} \}$ and its gap $\Delta_{(l) \tilde{\Theta}} = S_{\theta^* \tilde{\Theta}} - S_{(l) \tilde{\Theta}}$
Algorithm related	
CSE	The generic <i>combinatorial successive elimination</i> algorithm (Algorithm 1)
f_ρ	Function from $[k]$ to $[k]$, $f_\rho(x) = \lfloor x/2^\rho \rfloor$ for a $\rho \in (0, \log_2(k)]$ specifying the nature of the configuration elimination strategy
n_0	Variable to specify the initial sample size, $n_0 \in (N_{\alpha, \delta}, 2N_{\alpha, \delta}]$
n_e	$n_e = \lceil n_0/2^e \rceil + 1$ number of considered configurations in epoch $e \in [E]$
E	Number of epochs $E = \lceil \log_2(n_0/n_0 - N_{\alpha, \delta}) \rceil$
C_1	Internal constant variable for AC-Band, $C_1 = \log_{1+\frac{k-1}{E}}(2)$
C_2	Internal constant variable for AC-Band, $C_2 = 1 + \log_{1+\frac{k-1}{E}}\left(n_0 + 4\frac{n_0}{n_0 - N_{\alpha, \delta}}\right)$
C_3	Internal constant variable for AC-Band, $C_3 = \left\lceil \log_{1+\frac{k-1}{E}}(k) \right\rceil$
c_e	Internal variable for AC-Band in epoch $e \in [E]$: $c_e = \frac{(C_1 E - (2^E - 1)(2C_1 - C_2 - C_3))2^e}{2^E(-eC_1 + C_2 + C_3)}$
B_e	Budget for call of CSE in epoch $e \in [E]$: $B_e = B/c_e$ for an overall budget of B for AC-Band
R^{ρ_e, k, n_e}	Number of rounds in CSE in epoch $e \in [E]$
$P_r^{\rho_e, k, n_e}$	Number of partitions in CSE in epoch $e \in [E]$ and round $r \in [R^{\rho_e, k, n_e}]$
$\mathbb{A}_r(\theta)$	The partition in round r of CSE containing θ (emptyset otherwise)
b_r	Budget used in round r of CSE for a partition

B Extension of CSE for Finding ϵ -best Arms

B.1 Adjustments of Algorithm Parameters

We modify the definition of the function f in (Brandt et al. 2022) and thus define $f_\rho(x) = \lfloor \frac{x}{2^\rho} \rfloor$ for a $\rho \in (0, \log_2(k)]$ and obtain for

- $\rho = 1$: Combinatorial Successive Halving (CSH) with $f_\rho(k) = \lfloor \frac{k}{2} \rfloor$
- $\rho = \log_2(k)$: Combinatorial Successive Winner Stays (CSWS) with $f_\rho(k) = 1$
- $\rho \rightarrow 0$: Combinatorial Successive Rejects (CSR) with $f_\rho(k) = k - 1$.

Note that for a fixed subset size k and for a fixed $\rho \in (0, \log_2(k)]$, one can derive the number of rounds and number of partitions per round as follows. The number of rounds in the first while loop of Algorithm 1 can be computed as

$$R_1^{\rho,k,n} = \{ \min x : g^{(\circ x)}(n) \leq k \} \text{ for } g(x) = f_\rho(k) \cdot \left\lfloor \frac{x}{k} \right\rfloor + x \pmod{k},$$

where $g^{(\circ x)}$ denotes the x -times composition of g . Furthermore, it holds that $R_1^{\rho,k,n} \leq \lceil \log_{\frac{k}{f_\rho(k)}}(n) \rceil$, which we use for the sake of ease to estimate $R_1^{\rho,k,n}$ in the following in our theoretical analyses. In the second while-loop of Algorithm 1, we have only k arms left, thus we can calculate the number of rounds as

$$R_2^{\rho,k} = \left\{ \min x : f_\rho^{(\circ x)}(k) \leq 1 \right\}.$$

For all $k \in \mathbb{N}$, $R_2^{\rho,k} \leq \lceil \log_{\frac{k}{f_\rho(k)}}(k) \rceil$. The overall number of rounds is therefore $R^{\rho,k,n} = R_1^{\rho,k,n} + R_2^{\rho,k} \leq \lceil \log_{\frac{k}{f_\rho(k)}}(n) \rceil + \lceil \log_{\frac{k}{f_\rho(k)}}(k) \rceil$. The number of partitions per round $r \in \{1, \dots, R^{\rho,k,n}\}$ is given by

$$P_r^{\rho,k,n} = \left\lfloor \frac{n}{k} \left(\frac{f_\rho(k)}{k} \right)^{r-1} \right\rfloor.$$

Thus, we can automatically compute the number of rounds and partitions in Algorithm 1 if the parameter ρ for the discard function f_ρ , the subset size k , and the number of arms n are given. In contrast to (Brandt et al. 2022), we do not need to estimate and specify R and $\{P_r\}$ by hand before we can run Algorithm 1.

B.2 Theoretical Guarantees

The first step for extending Algorithm 1 to the context of AC is to extend the theoretical guarantees to the context of finding an ϵ -best arm of a finite set of arms (configurations in our terminology) and to derive a sufficient budget for CSE that is necessary to find such an ϵ -best arm, provided that an ϵ -best arm is defined as follows.

Definition B.1. Assume we have n arms (configurations) $\theta_1, \dots, \theta_n$. Let θ^* be such that

$$\forall \theta^* \in \Theta_{[2,k]}(\theta^*) \quad S_{\theta^*|\bar{\Theta}} \geq S_{(1)|\bar{\Theta}} - \epsilon.$$

We call θ^* an ϵ -best arm.

Note that we only have a finite set of n arms (configurations) $\theta_1, \dots, \theta_n$, which we identify simply by their indices $1, \dots, n$ in the following. Further assume in the following that an ϵ -best arm exists. We identify this arm by i^* , and write $\Delta_{i|\bar{\Theta}} = S_{(1)|\bar{\Theta}} - S_{i|\bar{\Theta}}$. In addition, let \mathbb{A}_r be the set of active arms in round $r \in [R^{\rho,k,n}]$, which will then be partitioned into $\mathbb{A}_{r,1}, \dots, \mathbb{A}_{r,P_r^{\rho,k,n}}$.

Theorem B.2 (Sufficient budget for CSE for finding an ϵ -best arm). Using Algorithm 1 with n arms, a discard function f_ρ , and a subset size k returns an arm i^* , which is an ϵ -best arm if the budget B is larger than

$$z(\rho, k, n, \epsilon) := R^{\rho,k,n} \max_{r \in [R^{\rho,k,n}]} P_r^{\rho,k,n} \cdot \max_{r \in [R^{\rho,k,n}]} \left(1 + \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho,k,n}]} \frac{\Delta_{(f_\rho(|\mathbb{A}_r(i^*)|)+1)|\mathbb{A}_r(i^*)}}{2} \right\} \right) \right),$$

where $\mathbb{A}_r(i^*)$ denotes the partition in round r of CSE containing i^* (or θ^*).

Proof of Theorem B.2. Step 1: Algorithm 1 never requires a number of problem instances that exceeds the budget B :

$$\sum_{r=1}^{R^{\rho,k,n}} P_r^{\rho,k,n} \cdot b_r = \sum_{r=1}^{R^{\rho,k,n}} P_r^{\rho,k,n} \cdot \left\lfloor \frac{B}{P_r^{\rho,k,n} \cdot R^{\rho,k,n}} \right\rfloor \leq \sum_{r=1}^{R^{\rho,k,n}} \frac{B}{R^{\rho,k,n}} = B.$$

Step 2: Assume in the following $B \geq z(\rho, k, n, \epsilon)$, then we have for each round $r \in [R^{\rho, k, n}]$ that

$$\begin{aligned} b_r &\geq \frac{B}{P_r^{\rho, k, n} \cdot R^{\rho, k, n}} - 1 \\ &\geq \max_{r \in [R^{\rho, k, n}]} \left(1 + \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{(f_\rho(|\mathbb{A}_r(i^*)|)+1)|\mathbb{A}_r(i^*)}}{2} \right\} \right) \right) - 1 \\ &= \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{(f_\rho(|\mathbb{A}_r(i^*)|)+1)|\mathbb{A}_r(i^*)}}{2} \right\} \right). \end{aligned}$$

We can assume in the following w.l.o.g. $i^* = 1$ and $\mathbb{A}_r(1) = \mathbb{A}_{r1}$ by relabeling the arms (configurations) and query sets (sets of configurations). Now, we first show, that $s_{1|\mathbb{A}_{r1}}(t) - s_{i|\mathbb{A}_{r1}}(t) \geq 0$ for all rounds $r \in [R^{\rho, k, n}]$, all arms $i \in \mathbb{A}_{r1}$ and all $t \geq \tau_i := \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{i|\mathbb{A}_{r1}}}{2} \right)$. Define $\gamma_{\theta|\bar{\Theta}}$ as the pointwise minimal function such that $|s_{\theta|\bar{\Theta}}(t) - S_{\theta|\bar{\Theta}}| \leq \gamma_{\theta|\bar{\Theta}}(t)$ for all t and $\gamma_{\bar{\Theta}} = \max_{\theta \in \bar{\Theta}} \gamma_{\theta|\bar{\Theta}}$. Thus, $\tau_i \geq \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\frac{\Delta_{i|\mathbb{A}_{r1}}}{2} \right)$ for all $r \in [R^{\rho, k, n}]$, and according to the definition of $\gamma_{\bar{\Theta}}$ we have

$$|s_{i|\mathbb{A}_{r1}}(t) - S_{i|\mathbb{A}_{r1}}| \leq \gamma_{\mathbb{A}_{r1}}(t) \leq \frac{\Delta_{i|\mathbb{A}_{r1}}}{2} \quad \text{for } t \geq \tau_i.$$

Thus, for all $t \geq \tau_i$,

$$\begin{aligned} s_{1|\mathbb{A}_{r1}}(t) - s_{i|\mathbb{A}_{r1}}(t) &= s_{1|\mathbb{A}_{r1}}(t) - S_{1|\mathbb{A}_{r1}} + S_{1|\mathbb{A}_{r1}} - S_{i|\mathbb{A}_{r1}} + S_{i|\mathbb{A}_{r1}} - s_{i|\mathbb{A}_{r1}}(t) \\ &= s_{1|\mathbb{A}_{r1}}(t) - S_{1|\mathbb{A}_{r1}} - (s_{i|\mathbb{A}_{r1}}(t) - S_{i|\mathbb{A}_{r1}}) + S_{1|\mathbb{A}_{r1}} - S_{i|\mathbb{A}_{r1}} \\ &\geq -2\gamma_{\mathbb{A}_{r1}}(t) + S_{1|\mathbb{A}_{r1}} - S_{i|\mathbb{A}_{r1}} \\ &\geq -2\frac{S_{1|\mathbb{A}_{r1}} - S_{i|\mathbb{A}_{r1}}}{2} + S_{1|\mathbb{A}_{r1}} - S_{i|\mathbb{A}_{r1}} = 0. \end{aligned}$$

In this scenario, arm i will be eliminated before arm 1, since the $f_\rho(|\mathbb{A}_r|)$ arms with the lowest statistic s are discarded in each round $r \in [R^{\rho, k, n}]$.

Now consider a round $r \in [R^{\rho, k, n}]$ and assume that by reordering the arms, w.l.o.g., $S_{1|\mathbb{A}_{r1}} \geq S_{2|\mathbb{A}_{r1}} \geq \dots \geq S_{|\mathbb{A}_{r1}||\mathbb{A}_{r1}}$. Since $S_{i|\mathbb{A}_{r1}}$ is non-increasing in i , the τ_i are non-increasing in i : $\tau_2 \geq \tau_3 \geq \dots$. Thus,

$$t \geq \tau_i \Rightarrow s_{1|\mathbb{A}_{r1}}(t) \geq s_{i|\mathbb{A}_{r1}}(t). \quad (2)$$

Assume now that $1 \in \mathbb{A}_r \wedge 1 \notin \mathbb{A}_{r+1}$

$$\begin{aligned} &\Rightarrow \sum_{i \in \mathbb{A}_{r1}} \mathbf{1}\{s_{i|\mathbb{A}_{r1}}(b_r) > s_{1|\mathbb{A}_{r1}}(b_r)\} > f_\rho(|\mathbb{A}_{r1}|) \\ &\Rightarrow \sum_{i \in \mathbb{A}_{r1}} \mathbf{1}\{b_r < \tau_i\} > f_\rho(|\mathbb{A}_{r1}|) \\ &\Rightarrow b_r < \tau_{f_\rho(|\mathbb{A}_{r1}|)+1}. \end{aligned}$$

This implies

$$1 \in \mathbb{A}_r \wedge b_r \geq \tau_{f_\rho(|\mathbb{A}_{r1}|)+1} \Rightarrow 1 \in \mathbb{A}_{r+1}. \quad (3)$$

Recall that $b_r \geq \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{(f_\rho(|\mathbb{A}_r(i^*)|)+1)|\mathbb{A}_r(i^*)}}{2} \right\} \right)$ and $\tau_{f(|\mathbb{A}_{r1}|)+1} = \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{f_\rho(|\mathbb{A}_{r1}|)+1|\mathbb{A}_{r1}}}{2} \right)$.

Case 1: $\max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{f_\rho(|\mathbb{A}_{r1}|)+1|\mathbb{A}_{r1}}}{2} \geq \frac{\epsilon}{2}$ and $1 \in \mathbb{A}_r$.

We have $b_r \geq \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{f_\rho(|\mathbb{A}_{r1}|)+1|\mathbb{A}_{r1}}}{2} \right) = \tau_{f_\rho(|\mathbb{A}_{r1}|)+1}$. By (3) we have that $1 \in \mathbb{A}_{r+1}$.

Case 2: $\max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{f_\rho(|\mathbb{A}_{r1}|)+1|\mathbb{A}_{r1}}}{2} < \frac{\epsilon}{2}$ and $1 \in \mathbb{A}_r$.

We have $b_r \geq \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\frac{\epsilon}{2} \right)$ and

$$\max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\frac{\epsilon}{2} \right) < \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{f_\rho(|\mathbb{A}_{r1}|)+1|\mathbb{A}_{r1}}}{2} \right) = \tau_{f_\rho(|\mathbb{A}_{r1}|)+1}.$$

If $1 \in \mathbb{A}_{r+1}$, Algorithm 1 continues in round $r+1$ with case 1 or case 2. Now assume $1 \notin \mathbb{A}_{r+1}$ and let

$$p := \min \left\{ i \in \mathbb{A}_{r1} : \max_{r \in [R^{\rho, k, n}]} \frac{S_{1|\mathbb{A}_{r1}} - S_{i|\mathbb{A}_{r1}}}{2} \geq \frac{\epsilon}{2} \right\}.$$

By the assumption of the case 2, we have $p > f_\rho(|\mathbb{A}_{r1}|) + 1$ and

$$b_r \geq \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\frac{\epsilon}{2} \right) \geq \max_{r \in [R^{\rho, k, n}]} \bar{\gamma}_{\mathbb{A}_{r1}}^{-1} \left(\max_{r \in [R^{\rho, k, n}]} \frac{\Delta_{i|\mathbb{A}_{r1}}}{2} \right) = \tau_i \text{ for all } i \geq p.$$

Thus, by (2) we have $s_{1|\mathbb{A}_{r1}}(b_r) \geq s_{i|\mathbb{A}_{r1}}(b_r)$ for $i \geq p$, so all arms $i \geq p$ are eliminated before arm 1. Moreover, we have

$$\max_{i \in \mathbb{A}_{r+1}} S_{i|\mathbb{A}_{r1}} \geq \max_{i < p} S_{i|\mathbb{A}_{r1}} \geq S_{1|\mathbb{A}_{r1}} - \epsilon$$

since

$$\frac{S_{1|\mathbb{A}_{r1}} - S_{i|\mathbb{A}_{r1}}}{2} < \max_{r \in [R^{\rho, k, n}]} \frac{S_{1|\mathbb{A}_{r1}} - S_{i|\mathbb{A}_{r1}}}{2} < \frac{\epsilon}{2}.$$

In addition, by definition of p it holds even for all $r \in [R^{\rho, k, n}]$ that $S_{i|\mathbb{A}_{r1}} > S_{1|\mathbb{A}_{r1}} - \epsilon$. Thus, although CSE might discard i^* , it is ensured that an ϵ -best arm is still active. Consequently, we relabel the arms such that the still active ϵ -best arm is now denoted by i^* .

Case 3: $1 \notin \mathbb{A}_r$.

Since $1 \in \mathbb{A}_0$, there was a round $\tilde{r} < r$ such that $1 \in \mathbb{A}_{\tilde{r}}$ and $1 \notin \mathbb{A}_{\tilde{r}+1}$. For this round \tilde{r} only case 2 was possible, otherwise $1 \in \mathbb{A}_{\tilde{r}+1}$. Since case 2 was true and $1 \notin \mathbb{A}_{\tilde{r}+1}$, we have

$$\max_{i \in \mathbb{A}_{\tilde{r}+1}} S_{i|\mathbb{A}_{\tilde{r}}} \geq S_{1|\mathbb{A}_{\tilde{r}}} - \epsilon$$

and also for all other rounds $r \in [R^{\rho, k, n}]$ that $S_{i|\mathbb{A}_{r1}} > S_{1|\mathbb{A}_{r1}} - \epsilon$.

□

C Guarantees for AC-Band

C.1 Number of Epochs

Let $N_{\alpha, \delta} = \lceil \frac{\ln(\delta)}{\ln(1-\alpha)} \rceil$ be the number of sampled configurations necessary to ensure that an ϵ -best configuration is contained in the samples with probability at least $1 - \delta$ provided the proportion of ϵ -best configurations in the configuration space is α (see Assumption (A2)). In each epoch e , we consider in total $n_e = \lceil \frac{n_0}{2^e} \rceil + 1$ configurations and keep the winner of the last epoch, such that we have to sample at least $\lceil \frac{n_0}{2^e} \rceil$ new configurations in epoch e . To be precise, we sample in one run of AC-Band, the following number of configurations in total:

$$\begin{aligned} \sum_{e=1}^E \left\lceil \frac{n_0}{2^e} \right\rceil &\geq \sum_{e=1}^E \frac{n_0}{2^e} \\ &= n_0 \sum_{e=1}^E \left(\frac{1}{2} \right)^e \\ &= n_0 \left(\sum_{e=0}^E \left(\frac{1}{2} \right)^e - 1 \right) \\ &= 2n_0 \left(1 - \frac{1}{2^{E+1}} \right) - n_0 \\ &= n_0 - \frac{n_0}{2^E}. \end{aligned}$$

This value must be greater than $N_{\alpha, \delta}$ to guarantee that we have an ϵ -best configuration contained in the sampled configurations with probability at least $1 - \delta$. Rearranging the above inequality leads to

$$\begin{aligned} 2^E &\geq \frac{n_0}{n_0 - N_{\alpha, \delta}} \\ \Rightarrow E &\geq \log_2 \left(\frac{n_0}{n_0 - N_{\alpha, \delta}} \right), \end{aligned}$$

where we consider that $n_0 > N_{\alpha,\delta}$. Since we want our algorithm to finish as fast as possible, $E = \left\lceil \log_2 \left(\frac{n_0}{n_0 - N_{\alpha,\delta}} \right) \right\rceil$ is a suitable choice. In addition, we need to guarantee that the number of epochs is well defined, thus we must ensure that

$$\begin{aligned} E &\geq \log_2 \left(\frac{n_0}{n_0 - N_{\alpha,\delta}} \right) \stackrel{!}{\geq} 1 \\ &\Leftrightarrow \frac{n_0}{n_0 - N_{\alpha,\delta}} \geq 2 \\ &\Leftrightarrow n_0 \leq 2N_{\alpha,\delta}. \end{aligned}$$

Putting everything together, we get the condition that $n_0 \in \left(N_{\alpha,\delta}, 2N_{\alpha,\delta} \right]$.

C.2 Sufficient Budget

Lemma C.1. *For $N \in \mathbb{N}$ and any $a, b, c \in \mathbb{R}$ it holds that*

$$\sum_{i=1}^N \frac{-ia + b + c}{2^i} = \frac{aN - (2^N - 1)(2a - b - c)}{2^N}.$$

Proof.

$$\begin{aligned} \sum_{i=1}^N \frac{-ia + b + c}{2^i} &= -a \cdot \sum_{i=1}^N \frac{i}{2^i} + (b + c) \cdot \sum_{i=1}^N \frac{1}{2^i} \\ &= -a \cdot \sum_{i=0}^N \frac{i}{2^i} + (b + c) \cdot \left(\sum_{i=0}^N \frac{1}{2^i} - 1 \right) \\ &= -a \frac{\frac{N}{2^{N+2}} - \frac{N+1}{2^{N+1}} + \frac{1}{2}}{\left(\frac{1}{2} - 1\right)^2} + (b + c) \left(\frac{1 - \frac{1}{2^{N+1}}}{1 - \frac{1}{2}} - 1 \right) \\ &= \frac{-aN}{2^N} + \frac{a(N+1)}{2^{N-1}} - 2a + (b + c) \left(1 - \frac{1}{2^N} \right) \\ &= \frac{-aN + a(2N+2) - a2^{N+1}}{2^N} + \frac{(b+c)(2^N - 1)}{2^N} \\ &= \frac{aN - (2^N - 1)(2a - b - c)}{2^N}, \end{aligned}$$

where the closed-form sum formulas of the geometric series are used in the third equality. \square

Proof of Thm. 5.1. If the total budget B for AC-Band is such that the epoch-wise budget for CSE is at least $z(\rho_e, k, n_e, \epsilon)$ for each epoch e (see Theorem B.2), AC-Band will return a configuration that is locally ϵ -best. With the help of Assumption (A3), we can then infer the claim.

Thus, the sufficient budget to guarantee that AC-Band finds a local ϵ -best configuration i^* can be computed as

$$\begin{aligned} &\sum_{e=1}^E z(\rho_e, k, n_e, \epsilon) \\ &= \sum_{e=1}^E R^{\rho_e, k, n_e} \max_{r \in [R^{\rho_e, k, n_e}]} P_r^{\rho_e, k, n_e} \cdot \max_{r \in [R^{\rho_e, k, n_e}]} \left(1 + \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho_e, k, n_e}]} \frac{\Delta(f_\rho(|\mathbb{A}_r(i^*)|+1)|_{\mathbb{A}_r(i^*)})}{2} \right\} \right) \right) \\ &= \sum_{e=1}^E (R_1^{\rho_e, k, n_e} + R_2^{\rho_e, k}) \left\lfloor \frac{n_e}{k} \right\rfloor \cdot \max_{r \in [R^{\rho_e, k, n_e}]} \left(1 + \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho_e, k, n_e}]} \frac{\Delta(f_\rho(|\mathbb{A}_r(i^*)|+1)|_{\mathbb{A}_r(i^*)})}{2} \right\} \right) \right) \\ &=: (*) \end{aligned}$$

Note that for $\rho_e = \log_2 \left(\frac{e+k-1}{e} \right)$, we have $f_{\rho_e}(x) = \lfloor \frac{x}{2^{\rho_e}} \rfloor = \lfloor \frac{xe}{e+k-1} \rfloor$. We can estimate $R_1^{\rho_e, k, n_e}$ now by

$$R_1^{\rho_e, k, n_e} \leq \left\lceil \log_{\frac{k}{f_{\rho_e}(k)}}(n_e) \right\rceil$$

$$\begin{aligned}
&= \left\lceil \log_{\frac{k}{\lfloor \frac{ke}{e+k-1} \rfloor}} \left(\left\lceil \frac{n_0}{2^e} \right\rceil + 1 \right) \right\rceil \\
&\leq \left\lceil \log_{\frac{e+k-1}{e}} \left(\left\lceil \frac{n_0}{2^e} \right\rceil + 1 \right) \right\rceil \\
&= \left\lceil \frac{\log \left(\left\lceil \frac{n_0}{2^e} \right\rceil + 1 \right)}{\log \left(1 + \frac{k-1}{e} \right)} \right\rceil \\
&\leq \left\lceil \frac{\log \left(\frac{n_0}{2^e} + 2 \right)}{\log \left(1 + \frac{k-1}{E} \right)} \right\rceil \\
&= \left\lceil \frac{\log (n_0 + 2^{e+1}) - \log(2^e)}{\log \left(1 + \frac{k-1}{E} \right)} \right\rceil \\
&\leq \left\lceil \frac{\log (n_0 + 2^{E+1}) - e \log(2)}{\log \left(1 + \frac{k-1}{E} \right)} \right\rceil \\
&\leq \left\lceil \frac{\log \left(n_0 + 2^{\log_2 \left(\frac{n_0}{n_0 - N_{\alpha, \delta}} \right) + 2} \right) - e \log(2)}{\log \left(1 + \frac{k-1}{E} \right)} \right\rceil \\
&= \left\lceil \frac{\log \left(n_0 + 4 \frac{n_0}{n_0 - N_{\alpha, \delta}} \right)}{\log \left(1 + \frac{k-1}{E} \right)} - e \cdot \frac{\log(2)}{\log \left(1 + \frac{k-1}{E} \right)} \right\rceil \\
&\leq \underbrace{1 + \log_{1 + \frac{k-1}{E}} \left(n_0 + 4 \frac{n_0}{n_0 - N_{\alpha, \delta}} \right)}_{=: C_2} - \underbrace{e \cdot \log_{1 + \frac{k-1}{E}}(2)}_{=: C_1}.
\end{aligned}$$

We can proceed analogously to get an upper bound for $R_2^{\rho_e, k}$:

$$\begin{aligned}
R_2^{\rho_e, k} &\leq \left\lceil \log_{\frac{k}{f_{\rho_e}(k)}} (k) \right\rceil = \left\lceil \log_{\frac{k}{\lfloor \frac{ke}{e+k-1} \rfloor}} (k) \right\rceil \\
&\leq \left\lceil \log_{\frac{e+k-1}{e}} (k) \right\rceil = \left\lceil \log_{1 + \frac{k-1}{e}} (k) \right\rceil \\
&\leq \left\lceil \log_{1 + \frac{k-1}{E}} (k) \right\rceil =: C_3 \leq C_2.
\end{aligned}$$

In the next step we can put the above estimations together.

$$\begin{aligned}
(*) &\leq \sum_{e=1}^E \left[\left(\frac{n_0}{2^e} + 2 \right) \frac{1}{k} \right] \cdot (-eC_1 + C_2 + C_3) \\
&\quad \cdot \max_{r \in [R^{\rho_e, k, n_e}]} \left(1 + \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho_e, k, n_e}]} \frac{\Delta(f_\rho(|\mathbb{A}_r(i^*)|+1)|_{\mathbb{A}_r(i^*)})}{2} \right\} \right) \right) \\
&\leq \underbrace{\max_{e \in [E]} \max_{r \in [R^{\rho_e, k, n_e}]} \left(1 + \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho_e, k, n_e}]} \frac{\Delta(f_\rho(|\mathbb{A}_r(i^*)|+1)|_{\mathbb{A}_r(i^*)})}{2} \right\} \right) \right)}_{=: \bar{\gamma}^{-1}} \\
&\quad \cdot \left(\frac{n_0}{k} \sum_{e=1}^E \frac{-eC_1 + C_2 + C_3}{2^e} - \frac{2C_1}{k} \sum_{e=1}^E e + \frac{2E(C_2 + C_3)}{k} \right) \\
&= \bar{\gamma}^{-1} \cdot \left(\frac{n_0}{k} \cdot \frac{C_1 E - (2^E - 1)(2C_1 - C_2 - C_3)}{2^E} + \frac{2E(C_2 + C_3) - C_1 E(E + 1)}{k} \right)
\end{aligned}$$

according to Lemma C.1 and the Gaussian sum formula.

A cruder bound for the sufficient budget is given by

$$\bar{\gamma}^{-1} \cdot \frac{n_0 + 2^{E+1}}{k} \cdot \frac{C_1 E - (2^E - 1)(2C_1 - C_2 - C_3)}{2^E}. \tag{4}$$

Recall that the number of parallel runs is decreasing with the epoch e , so that in the worst case the configuration returned by AC-Band is sampled only in the last epoch E . Consequently it will be run in parallel only with configurations, which are considered in the last epoch E , and guaranteed to be a local ϵ -best configuration (see Theorem B.2). By Assumption (A3), the local ϵ -best property corresponds to a global ϵ -best property with probability of at least

$$1 - \frac{1}{\#\text{query sets containing } i^* \text{ in epoch } E} = 1 - (R^{\rho_E, k, n_E})^{-1}$$

for this worst case. In addition we have to take into account that an ϵ -best configuration is contained only with probability at least $1 - \delta$, such that we get an overall probability of at least

$$\min\{1 - \delta, 1 - (R^{\rho_E, k, n_E})^{-1}\}$$

that AC-Band returns an ϵ -best configuration. \square

Note that the total budget of AC band is divided among the epoch-wise calls of CSE by means of the quotient c_e :

$$c_e = \frac{(C_1 E - (2^E - 1)(2C_1 - C_2 - C_3))2^e}{2^E(-eC_1 + C_2 + C_3)}.$$

This quotient is obtained by bounding the sufficient budget for CSE similar as in the proof of Theorem 5.1 to obtain (4):

$$\begin{aligned} z(\rho_e, k, n_e, \epsilon) &= R^{\rho_e, k, n_e} \max_{r \in [R^{\rho_e, k, n_e}]} P_r^{\rho_e, k, n_e} \\ &\cdot \max_{r \in [R^{\rho_e, k, n_e}]} \left(1 + \bar{\gamma}_{\mathbb{A}_r(i^*)}^{-1} \left(\max \left\{ \frac{\epsilon}{2}, \max_{r \in [R^{\rho_e, k, n_e}]} \frac{\Delta(f_\rho(|\mathbb{A}_r(i^*)|+1)|_{\mathbb{A}_r(i^*)})}{2} \right\} \right) \right) \\ &\leq \bar{\gamma}^{-1} (R_1^{\rho_e, k, n_e} + R_2^{\rho_e, k}) \left\lfloor \frac{n_e}{k} \right\rfloor \\ &\leq \bar{\gamma}^{-1} \cdot \frac{n_0 + 2^{E+1}}{k} \cdot \frac{-eC_1 + C_2 + C_3}{2^e}. \end{aligned}$$

AC-Band thus allocates its entire budget such that any call to CSE is guaranteed to return an appropriate configuration once the total budget is sufficiently large.

C.3 Discussion of Sufficient Budget

The behavior of the sufficient budget with respect to k , α and δ is illustrated in Figure 3. Note that we ignore the γ^{-1} terms in these plots, as these depend on the underlying AC problem and occur as a multiplicative constant in the sufficient budget.




Figure 3: Sufficient budget for different values for k , α and δ .

The dependency of the sufficient budget on k , α and δ is as expected, since it decreases with increasing k , α and δ , respectively.

D Extended Experiments

We provide further details regarding the experiments in Section 6. In particular, the results from Figure 2 are reported in more detail in Table 1 and augmented by the results for the Regions200 dataset, for which we provide a similar illustration of the results in Figure 4. Moreover, we outline additional experimental results and provide additional metrics to evaluate the quality of the configurations found. The experimental setup used here is the same as in the main paper. We report two additional metrics to

provide additional insights: the percent gap to subset-best and the R^δ metric used in previous works (Kleinberg et al. 2019; Weisz, György, and Szepesvári 2018, 2019; Weisz et al. 2020) within the following tables. The percent gap to subset-best is a variation of the percent gap to best metric where only the configurations that were sampled by the method during a run are considered. In this way, we can see how good a method performs within the sample it selects. A value of 0 means the configuration returned is the best in the subset. A 10% cutoff is used for the δ -capped runtime. We provide results for two additional experiments: (i) varying the δ for AC-Band and (I)CAR(++), and, (ii) increasing the configuration sampling budget of AC-band.




Figure 4: Mean CPU time and percent gap to best over 5 seeds for $\delta = 0.05$ and different α (columns) for AC-Band, ICAR, CAR++ and Hyperband on the Regions200 dataset. Circles indicate variants of AC-Band. Rectangles represent the standard error over the seeds. The number of configurations tried for CAR++: {97, 245, 492}, ICAR: {134, 351, 724}, AC-Band: {60, 153, 303}, Hyperband($\eta = 5$): {842}, Hyperband($\eta = 8$): {618}.

Varying δ The user must decide δ based on their preferences, thus there is no “correct” value to set δ to in our experiments. Therefore, we also experiment with $\delta = 0.01$ in addition to the results in Section 6 where $\delta = 0.05$ is used. The results for a lower δ (see Figure 5 or Table 2) are consistent with the results reported for $\delta = 0.05$. In particular, AC-Band with $k = 2$ lies on the Pareto front of percent gap to best and CPU time, backing up the claim that a lower value of k is preferable. With $\delta = 0.01$ and $k = 2$, AC-Band is 80% percent faster than ICAR and 74% faster than Hyperband over all α and all datasets, while providing configurations that are only 7% and 6% worse in terms of the gap to the best configuration. Furthermore, AC-Band exhibits a nearly linear speedup with the number of available cores, leading to a low wallclock time for increasing k . We note that a real parallel implementation would, of course, suffer from some extra overhead.

Looking closer at the results obtained for the CNFuzzDD dataset for $\alpha = 0.01$, we note that the percent gap to best increases for $k = 8$. This increase is solely due to one seed for which a percent gap to best value of 5.74 is obtained. AC-Band uses most of its sampling budget in the first rounds, where large sets of configurations are evaluated on large sets of instances. Over the course of AC-Band, both the configuration and instance sets become smaller. For the seed in question, AC-Band samples one new configuration in the last round that is able to beat the current incumbent on 18 instances. Since the configuration wins, it is returned, even though the incumbent has seen more instances and proven worthy over the previous 8 epochs. The possibility of this happening grows with the number of configurations to sample (decreasing α) since the same amount of instances is distributed among more configurations.

Increasing the sampling budget of AC-Band AC-Band samples fewer configurations than Hyperband or (I)CAR(++), and leaves more of the configuration space unexplored. This explains, in part, why AC-Band usually has a worse percent gap to best than its competitors, but less runtime. To investigate this further, we let AC-Band sample the same number of configurations as CAR++ and Hyperband with $\eta = 8$ in Table 3 and Table 4. In particular, we set N to be equal to the number of samples of either method and set $n_0 = N + 1$. Note that we do not sample exactly the same amount of configurations due to the rounding operations within AC-Band.

On the CNFuzzDD and RCW dataset (see Table 3), a larger sampling budget for AC-Band leads to a gap to best that is nearly as good as those obtained by (I)CAR(++), while needing significantly less CPU time. The same can be seen for the Regions200 dataset, however with some exceptions where the gap to best increases. Specifically, for $\alpha = 0.05$ where only 101 configurations out of 2000 are examined. For a few seeds, no good configuration is contained in the 101 samples, leading to these results.

Table 4 report the results obtained for Hyperband with different values of η as well as the AC-Band results with a sampling budget of 618 configurations. For all three datasets, AC-Band comes closer to the results of Hyperband in terms of gap to best, while needing significantly less CPU time. This is especially true for the Regions200 and RCW datasets. On the CNFuzzDD dataset, AC-Band’s performance is weaker, which may be due to this dataset containing the smallest number of instances.

α	CPU Time (thousand days)						Percent gap to best						Percent gap to subset-best						R^δ					
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01	
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
Method	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
ICAR	100.64	12.74	242.74	14.96	467.32	25.08	0.02	0.04	0.01	0.01	0.02	0.03	0.01	0.02	0.01	0.01	0.02	0.03	5.0	0.1	4.9	0.1	4.9	0.1
CAR++	92.30	5.48	224.21	16.38	452.06	18.08	0.05	0.04	0.01	0.01	0.01	0.01	0.02	0.03	0.00	0.01	0.01	0.01	5.2	0.2	4.9	0.1	4.9	0.1
CAR	157.76	18.07	367.86	7.24	771.45	21.72	0.04	0.03	0.01	0.01	0.01	0.01	0.01	0.02	0.00	0.01	0.01	0.01	5.2	0.2	4.9	0.1	4.9	0.1
$k = 2$	13.05	2.04	11.76	0.90	11.02	0.56	0.14	0.17	0.16	0.19	0.05	0.08	0.09	0.12	0.14	0.19	0.05	0.08	5.7	0.90	5.7	1.0	5.1	0.5
$k = 4$	16.32	1.91	14.91	0.68	14.82	0.53	0.05	0.06	0.02	0.03	0.01	0.02	0.00	0.00	0.00	0.01	0.01	0.02	5.3	0.3	5.1	0.2	4.9	0.1
$k = 8$	20.09	1.92	21.58	1.24	21.21	0.91	0.12	0.14	0.01	0.01	0.01	0.01	0.06	0.11	0.00	0.01	0.01	0.01	5.6	0.8	5.0	0.1	4.9	0.1
$k = 16$	37.87	3.46	36.07	1.71	38.45	0.32	0.05	0.06	0.05	0.07	0.05	0.10	0.02	0.03	0.03	0.05	0.05	0.10	5.2	0.3	5.0	0.3	5.2	0.6

(a): CNFuzzDD dataset

α	CPU Time (days)						Percent gap to best						Percent gap to subset-best						R^δ					
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01	
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
ICAR	164.30	91.05	274.84	100.72	420.15	103.24	0.24	0.16	0.09	0.09	0.04	0.04	0.00	0.00	0.00	0.00	0.00	0.00	34.8	4.3	29.8	2.2	28.5	1.8
CAR++	229.29	19.93	566.99	28.21	1097.90	88.41	0.27	0.17	0.15	0.07	0.09	0.09	0.00	0.00	0.00	0.00	0.00	0.00	35.3	4.3	32.0	2.2	29.8	1.8
CAR	523.69	53.34	1294.87	64.11	2549.22	199.00	0.27	0.17	0.16	0.09	0.09	0.09	0.00	0.00	0.10	0.30	0.00	0.00	35.3	4.5	31.9	1.6	29.8	2.2
$k = 2$	154.27	25.29	137.79	7.59	132.00	7.78	0.31	0.21	0.23	0.12	0.08	0.11	0.00	0.00	0.01	0.02	0.00	0.00	36.6	5.9	34.6	5.0	30.2	1.6
$k = 4$	227.07	33.18	200.72	20.76	194.88	12.07	0.29	0.17	0.19	0.17	0.12	0.12	0.00	0.00	0.00	0.00	0.01	0.02	36.2	4.4	34.3	5.4	31.8	2.7
$k = 8$	271.02	60.26	292.56	31.26	272.56	21.64	0.29	0.17	0.21	0.14	0.13	0.13	0.00	0.00	0.01	0.02	0.01	0.02	36.2	6.3	33.5	3.5	31.1	2.2
$k = 16$	513.57	142.68	469.19	46.49	519.70	42.30	0.50	0.48	0.32	0.27	0.10	0.10	0.00	0.00	0.03	0.06	0.00	0.00	40.0	9.0	36.8	6.9	30.6	1.7

(b): Regions200 dataset

α	CPU Time (thousand days)						Percent gap to best						Percent gap to subset-best						R^δ					
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01	
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
ICAR	1.28	0.39	2.03	0.30	4.07	0.24	0.14	0.08	0.08	0.03	0.06	0.02	0.00	0.01	0.00	0.00	0.00	0.00	156.1	11.9	146.5	4.1	143.3	4.9
CAR++	1.73	0.37	3.64	0.18	7.52	0.13	0.17	0.09	0.10	0.03	0.06	0.02	0.00	0.01	0.00	0.00	0.00	0.00	162.1	11.9	149.1	4.1	143.3	4.9
CAR	3.30	0.50	7.59	0.19	15.65	0.25	0.17	0.09	0.10	0.03	0.06	0.02	0.00	0.01	0.00	0.00	0.00	0.00	160.1	13.3	149.1	4.7	143.3	4.9
$k = 2$	0.30	0.02	0.27	0.01	0.26	0.01	0.18	0.13	0.12	0.12	0.10	0.04	0.02	0.03	0.04	0.07	0.04	0.04	164.1	32.2	152.2	17.8	144.7	2.9
$k = 4$	0.52	0.04	0.47	0.03	0.45	0.02	0.21	0.08	0.19	0.22	0.07	0.08	0.03	0.03	0.08	0.16	0.02	0.03	169.8	18.5	165.2	36.5	139.8	8.4
$k = 8$	0.69	0.09	0.72	0.03	0.71	0.04	0.20	0.11	0.23	0.22	0.09	0.09	0.04	0.06	0.14	0.15	0.06	0.07	165.6	24.0	173.0	42.5	143.0	14.0
$k = 16$	1.42	0.29	1.31	0.06	1.43	0.08	0.21	0.08	0.15	0.16	0.12	0.06	0.05	0.07	0.03	0.06	0.05	0.04	167.6	15.5	158.2	30.5	146.4	4.6

(c): RCW dataset

Table 1: Mean (μ) and standard derivation (σ) for CPU time, percent gap to best, percent gap to subset-best and mean δ -capped runtime over 5 seeds for $\delta = \mathbf{0.05}$ and different α (columns) for AC-Band, ICAR, CAR++, CAR on the CNFuzzDD (top), Regions200 (middle) and RCW (bottom) dataset. The number of configurations tried for CAR(++): {97, 245, 492}, ICAR: {134, 351, 724}, AC-Band: {60, 153, 303}.




Figure 5: Mean CPU time and percent gap to best over 5 seeds for $\delta = 0.01$ and different α (columns) for AC-Band, ICAR, CAR++ and Hyperband on CNFuzzDD (top), Regions200 (middle) and RCW (bottom). Circles indicate variants of AC-Band. Rectangles represent the standard error over the seeds. The number of configurations tried for CAR++: {128, 325, 652}, ICAR: {166, 431, 884}, AC-Band: {93, 232, 462}, Hyperband($\eta = 5$): {842}, Hyperband($\eta = 8$): {618}.

α	CPU Time (days)						Percent gap to best						Percent gap to subset-best						R^δ					
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01	
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
Method	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
ICAR	130.57	13.14	325.61	11.80	619.16	15.58	0.02	0.04	0.01	0.01	0.02	0.03	0.01	0.02	0.01	0.01	0.02	0.03	5.0	0.2	4.9	0.1	4.9	0.1
CAR++	128.14	15.59	327.18	9.35	654.46	12.68	0.02	0.04	0.01	0.01	0.03	0.02	0.01	0.02	0.01	0.01	0.03	0.02	5.0	0.2	4.9	0.1	4.9	0.1
CAR	220.33	29.17	554.95	19.47	1169.06	12.44	0.01	0.02	0.01	0.01	0.02	0.02	0.00	0.00	0.01	0.01	0.01	0.02	5.1	0.2	4.9	0.1	4.9	0.1
$k = 2$	11.56	0.77	11.17	0.32	11.20	0.65	0.10	0.14	0.00	0.00	0.08	0.15	0.08	0.15	0.00	0.00	0.08	0.14	5.5	0.8	4.9	0.1	5.4	0.8
$k = 4$	14.38	0.63	13.29	0.50	14.11	0.36	0.02	0.03	0.01	0.02	0.09	0.14	0.00	0.00	0.01	0.01	0.09	0.14	5.1	0.1	4.9	0.0	5.3	0.8
$k = 8$	21.43	2.45	20.38	0.63	21.98	0.52	0.02	0.03	0.10	0.14	2.16	2.29	0.00	0.00	0.09	0.12	1.16	1.29	5.1	0.1	5.4	0.7	13.0	16.1
$k = 16$	29.18	1.99	33.77	0.82	31.70	0.45	0.02	0.02	0.02	0.03	0.02	0.04	0.00	0.00	0.01	0.02	0.02	0.04	5.0	0.1	5.0	0.2	5.0	0.3

(a): CNFuzzDD dataset

α	CPU Time (days)						Percent gap to best						Percent gap to subset-best						R^δ					
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01	
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
Method	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
ICAR	226.12	127.05	374.51	131.49	579.64	145.09	0.24	0.16	0.09	0.09	0.04	0.04	0.01	0.03	0.00	0.00	0.00	0.00	34.8	4.3	29.8	2.2	28.5	1.8
CAR++	349.44	29.05	813.73	35.96	1604.52	133.59	0.24	0.16	0.11	0.08	0.04	0.04	0.00	0.00	0.00	0.00	0.00	0.00	34.8	4.3	30.6	2.2	28.5	1.8
CAR	798.6	74.21	1885.13	83.41	3717.16	264.71	0.24	0.16	0.11	0.08	0.04	0.04	0.00	0.00	0.00	0.00	0.00	0.00	34.8	4.3	30.6	1.6	28.5	1.8
$k = 2$	147.24	8.72	143.81	9.80	129.57	5.53	0.38	0.17	0.25	0.22	0.17	0.10	0.00	0.00	0.02	0.03	0.05	0.07	38.9	4.9	34.4	5.2	31.6	3.5
$k = 4$	186.61	17.15	182.01	17.80	172.98	7.71	0.40	0.15	0.31	0.23	0.17	0.10	0.00	0.00	0.03	0.03	0.01	0.02	39.6	3.5	36.0	5.6	31.6	3.5
$k = 8$	283.07	58.47	267.06	23.24	283.98	22.41	0.37	0.20	0.17	0.11	0.20	0.09	0.02	0.04	0.02	0.02	0.03	0.02	38.1	4.9	31.5	3.3	32.0	3.0
$k = 16$	368.87	43.25	433.53	28.73	403.93	27.27	0.26	0.17	0.18	0.12	0.17	0.11	0.00	0.00	0.00	0.00	0.02	0.02	35.2	5.4	33.0	2.8	31.2	3.0

(b): Regions200 dataset

α	CPU Time (thousand days)						Percent gap to best						Percent gap to subset-best						R^δ					
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01	
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
Method	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
ICAR	1.27	0.22	2.64	0.29	5.26	0.40	0.11	0.03	0.08	0.03	0.06	0.04	0.00	0.00	0.00	0.00	0.00	0.00	150.1	2.3	146.2	3.7	142.1	5.9
CAR++	2.20	0.32	5.21	0.25	10.82	0.38	0.14	0.08	0.08	0.03	0.06	0.02	0.00	0.00	0.00	0.00	0.00	0.00	156.1	2.3	146.5	3.7	143.3	5.9
CAR	4.44	0.48	10.79	0.26	22.60	0.61	0.14	0.08	0.08	0.03	0.06	0.02	0.00	0.00	0.00	0.00	0.00	0.00	156.1	11.9	146.5	4.1	143.3	4.9
$k = 2$	0.29	0.01	0.27	0.09	0.26	0.01	0.21	0.14	0.11	0.04	0.09	0.04	0.07	0.08	0.04	0.04	0.03	0.03	168.7	25.9	145.1	3.6	141.9	4.7
$k = 4$	0.45	0.03	0.42	0.01	0.42	0.02	0.16	0.12	0.10	0.04	0.09	0.05	0.05	0.09	0.04	0.04	0.04	0.04	156.9	19.6	144.9	3.3	141.8	4.6
$k = 8$	0.69	0.02	0.70	0.03	0.74	0.05	0.20	0.13	0.11	0.04	0.09	0.05	0.05	0.09	0.04	0.04	0.04	0.04	164.2	21.7	145.1	3.6	141.8	4.6
$k = 16$	1.05	0.04	1.23	0.05	1.17	0.05	0.18	0.15	0.12	0.05	0.07	0.06	0.06	0.09	0.04	0.04	0.02	0.02	163.7	27.2	145.6	3.8	139.4	5.0

(c): RCW dataset

Table 2: Mean (μ) and standard derivation (σ) for CPU time, percent gap to best, percent gap to subset-best and mean δ -capped runtime over 5 seeds for $\delta = \mathbf{0.01}$ and different values of α (columns) for AC-Band, ICAR, CAR++, CAR on the CNFuzzDD (top), Regions200 (middle) and RCW (bottom) datasets. The number of configurations tried for CAR(++): {128, 325, 652}, ICAR: {166, 431, 884}, AC-Band: {93, 232, 462}.

α	CPU Time (days)						Percent gap to best						Percent gap to subset-best						R^δ											
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01							
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ						
Method																														
ICAR	100.64	12.74	242.74	14.96	467.32	25.08	0.02	0.04	0.01	0.01	0.02	0.03	0.01	0.02	0.01	0.01	0.02	0.03	5.0	0.1	4.9	0.1	4.9	0.1	5.2	0.2	4.9	0.1	4.9	0.1
CAR++	92.30	5.48	224.21	16.38	452.06	18.08	0.05	0.04	0.01	0.01	0.01	0.01	0.02	0.03	0.00	0.01	0.01	0.01	5.2	0.2	4.9	0.1	4.9	0.1	5.2	0.2	4.9	0.1	4.9	0.1
CAR	157.76	18.07	367.86	7.24	771.45	21.72	0.04	0.03	0.01	0.01	0.01	0.01	0.01	0.02	0.00	0.01	0.01	0.01	5.2	0.2	4.9	0.1	4.9	0.1	5.2	0.2	4.9	0.1	4.9	0.1
$k = 2$	11.66	0.87	11.50	0.25	10.83	0.25	0.03	0.02	0.02	0.02	0.09	0.13	0.00	0.00	0.01	0.02	0.09	0.14	5.1	0.1	5.0	0.2	5.3	0.7	5.1	0.1	5.4	0.8	5.0	0.2
$k = 4$	14.18	0.49	14.53	0.42	13.51	0.35	0.03	0.03	0.10	0.16	0.03	0.04	0.00	0.00	0.09	0.14	0.03	0.05	5.1	0.1	5.4	0.8	5.0	0.2	5.1	0.1	5.4	0.8	5.0	0.2
$k = 8$	22.39	0.82	21.45	0.59	20.17	0.58	0.02	0.03	0.02	0.01	0.00	0.01	0.00	0.00	0.01	0.01	0.00	0.01	5.1	0.2	5.0	0.1	4.9	0.1	5.1	0.2	5.0	0.1	4.9	0.1
$k = 16$	29.56	1.61	30.00	1.18	32.28	0.43	0.04	0.03	0.01	0.01	0.02	0.03	0.00	0.00	0.00	0.00	0.02	0.03	5.2	0.2	5.0	0.1	5.0	0.1	5.2	0.2	5.0	0.1	5.0	0.1

(a): CNFuzzDD dataset

α	CPU Time (days)						Percent gap to best						Percent gap to subset-best						R^δ											
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01							
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ						
Method																														
ICAR	164.30	91.05	274.84	100.72	420.15	103.24	0.24	0.16	0.09	0.09	0.04	0.04	0.00	0.00	0.00	0.00	0.00	0.00	34.8	4.3	29.8	2.2	28.5	1.8	35.3	4.3	32.0	2.2	29.8	1.8
CAR++	229.29	19.93	566.99	28.21	1097.90	88.41	0.27	0.17	0.15	0.07	0.09	0.09	0.00	0.00	0.00	0.00	0.00	0.00	35.3	4.3	32.0	2.2	29.8	1.8	35.3	4.5	31.9	1.6	29.8	2.2
CAR	523.69	53.34	1294.87	64.11	2549.22	199.00	0.27	0.17	0.16	0.09	0.09	0.09	0.00	0.00	0.10	0.30	0.00	0.00	35.3	4.5	31.9	1.6	29.8	2.2	35.3	4.5	31.9	1.6	29.8	2.2
$k = 2$	148.29	10.60	136.75	12.34	125.29	3.91	0.35	0.20	0.15	0.13	0.04	0.04	0.01	0.03	0.01	0.02	0.00	0.01	37.8	5.6	31.8	3.0	28.9	1.8	37.8	5.6	31.8	3.0	28.9	1.8
$k = 4$	184.87	19.23	187.03	15.01	166.71	6.53	0.58	0.51	0.16	0.18	0.09	0.10	0.14	0.28	0.00	0.00	0.02	0.05	42.4	9.7	32.3	4.5	30.0	1.5	42.4	9.7	32.3	4.5	30.0	1.5
$k = 8$	310.18	25.52	263.16	19.92	241.98	13.16	0.39	0.27	0.12	0.09	0.14	0.14	0.04	0.08	0.01	0.01	0.05	0.10	38.6	6.5	30.7	1.6	31.1	2.4	38.6	6.5	30.7	1.6	31.1	2.4
$k = 16$	380.15	43.31	358.15	40.36	384.76	22.49	0.36	0.19	0.18	0.18	0.09	0.10	0.01	0.03	0.02	0.03	0.01	0.02	38.1	5.2	32.5	4.4	30.0	1.5	38.1	5.2	32.5	4.4	30.0	1.5

(b): Regions200 dataset

α	CPU Time (days)						Percent gap to best						Percent gap to subset-best						R^δ											
	0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01		0.05		0.02		0.01							
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ						
Method																														
ICAR	1.28	0.40	2.03	0.30	4.07	0.24	0.14	0.08	0.08	0.03	0.06	0.02	0.00	0.01	0.00	0.00	0.00	0.00	156.1	11.9	146.5	4.1	143.3	4.9	162.1	11.9	149.1	4.1	143.3	4.9
CAR++	1.72	0.37	3.64	0.18	7.52	0.13	0.17	0.09	0.10	0.03	0.06	0.02	0.00	0.01	0.00	0.00	0.00	0.00	162.1	11.9	149.1	4.1	143.3	4.9	162.1	11.9	149.1	4.1	143.3	4.9
CAR	3.30	0.50	7.59	0.19	15.65	0.26	0.17	0.09	0.10	0.03	0.06	0.02	0.00	0.01	0.00	0.00	0.00	0.00	160.1	13.3	149.1	4.7	143.3	4.9	160.1	13.3	149.1	4.7	143.3	4.9
$k = 2$	0.30	0.02	0.27	0.01	0.27	0.01	0.14	0.09	0.11	0.04	0.12	0.06	0.01	0.03	0.04	0.04	0.07	0.05	154.5	16.1	145.1	3.6	146.4	4.6	154.5	16.1	145.1	3.6	145.9	12.3
$k = 4$	0.43	0.02	0.04	0.01	0.04	0.01	0.18	0.11	0.11	0.04	0.10	0.10	0.04	0.03	0.03	0.04	0.07	0.06	161.3	24.1	145.1	3.6	145.9	12.3	161.3	24.1	145.1	3.6	145.9	12.3
$k = 8$	0.81	0.06	0.73	0.03	0.69	0.02	0.18	0.11	0.11	0.05	0.13	0.06	0.04	0.05	0.03	0.04	0.09	0.05	159.8	18.9	146.6	6.7	148.5	6.3	159.8	18.9	146.6	6.7	148.5	6.3
$k = 16$	1.09	0.10	1.06	0.04	1.17	0.04	0.23	0.20	0.09	0.05	0.15	0.06	0.08	0.11	0.03	0.04	0.10	0.06	169.8	39.2	141.8	4.6	151.8	6.3	169.8	39.2	141.8	4.6	151.8	6.3

(c): RCW dataset

Table 3: Mean (μ) and standard derivation (σ) for CPU time, percent gap to best, percent gap to subset-best and mean δ -capped runtime over 5 seeds for $\delta = \mathbf{0.05}$ and different values of α (columns) for AC-Band, ICAR, CAR++, CAR on the CNFuzzDD (top), Regions200 (middle) and RCW (bottom) datasets. For AC-Band N was set to the number of configurations sampled by CAR++ for the respective α value. The number of configurations tried for CAR(++): {97, 245, 492}, ICAR: {134, 351, 724}, AC-Band: {101, 247, 492}. Note that AC-Band samples slightly more configurations due to rounding operations.

Method		CPU Time (days)		Percent gap to best		Percent gap to subset-best		R^δ	
		μ	σ	μ	σ	μ	σ	μ	σ
Hyperband	$\eta = 4$	80.05	18.36	0.14	0.17	0.13	0.16	5.7	0.0
	$\eta = 5$	83.37	11.50	0.27	0.26	0.27	0.26	6.3	0.1
	$\eta = 6$	77.64	23.71	0.13	0.14	0.11	0.13	5.8	0.8
	$\eta = 8$	65.43	14.21	0.16	0.14	0.16	0.14	5.7	0.8
AC-Band	$k = 2$	10.55	0.32	0.39	0.53	0.39	0.53	7.1	3.0
	$k = 4$	13.66	0.26	0.06	0.08	0.06	0.08	5.2	0.5
	$k = 8$	19.49	0.26	0.21	0.36	0.21	0.36	6.1	2.1
	$k = 16$	33.35	0.57	0.01	0.02	0.01	0.02	5.0	0.1

(a): CNFuzzDD dataset

Method		CPU Time (days)		Percent gap to best		Percent gap to subset-best		R^δ	
		μ	σ	μ	σ	μ	σ	μ	σ
Hyperband	$\eta = 4$	798.58	64.32	0.06	0.06	0.06	0.05	29.0	2.3
	$\eta = 5$	668.78	120.01	0.10	0.11	0.04	0.05	30.7	3.1
	$\eta = 6$	662.96	94.83	0.06	0.07	0.06	0.06	28.6	2.4
	$\eta = 8$	559.19	64.26	0.07	0.08	0.00	0.00	29.9	2.5
AC-Band	$k = 2$	117.68	6.73	0.09	0.09	0.05	0.07	29.6	1.92
	$k = 4$	165.07	11.85	0.09	0.09	0.05	0.07	29.6	1.92
	$k = 8$	230.52	25.96	0.09	0.09	0.02	0.02	29.6	1.92
	$k = 16$	385.45	32.5	0.09	0.09	0.04	0.02	29.2	2.03

(b): Regions200 dataset

Method		CPU Time (days)		Percent gap to best		Percent gap to subset-best		R^δ	
		μ	σ	μ	σ	μ	σ	μ	σ
Hyperband	$\eta = 4$	991.58	55.48	0.07	0.08	0.07	0.07	141.4	9.1
	$\eta = 5$	883.04	67.03	0.12	0.08	0.07	0.07	147.6	12.2
	$\eta = 6$	873.19	56.11	0.10	0.11	0.10	0.10	144.3	22.1
	$\eta = 8$	796.60	59.49	0.04	0.08	0.00	0.00	138.0	9.3
AC-Band	$k = 2$	250.98	6.87	0.14	0.10	0.11	0.05	151.1	14.0
	$k = 4$	411.85	11.33	0.12	0.12	0.09	0.07	148.5	15.1
	$k = 8$	657.61	25.65	0.11	0.12	0.07	0.07	146.6	17.4
	$k = 16$	1196.41	38.75	0.08	0.07	0.07	0.07	139.3	16.0

(c): RCW dataset

Table 4: Mean (μ) and standard derivation (σ) for CPU time, percent gap to best, percent gap to subset-best and mean δ -capped runtime over 5 seeds for different η for Hyperband and AC-Band on the CNFuzzDD (top), Regions200 (middle) and RCW (bottom) datasets. The number of configurations tried for Hyperband: {378, 842, 280, 618}, AC-Band: {618}. Note that we use a value of η for which no more than the available number of configurations in a dataset would be sampled.