

UTILITARIAN ALGORITHM CONFIGURATION FOR INFINITE PARAMETER SPACES

Devon R. Graham & Kevin Leyton-Brown

Department of Computer Science
University of British Columbia
Vancouver, BC
{drgraham, kevinlb}@cs.ubc.ca

ABSTRACT

Utilitarian algorithm configuration is a general-purpose technique for automatically searching the parameter space of a given algorithm to optimize its performance, as measured by a given utility function, on a given set of inputs. Recently introduced utilitarian configuration procedures offer optimality guarantees about the returned parameterization while provably adapting to the hardness of the underlying problem. However, the applicability of these approaches is severely limited by the fact that they only search a finite, relatively small set of parameters. They cannot effectively search the configuration space of algorithms with continuous or uncountable parameters. In this paper we introduce a new procedure, which we dub COUP (Continuous, Optimistic Utilitarian Procrastination). COUP is designed to search infinite parameter spaces efficiently to find good configurations quickly. Furthermore, COUP maintains the theoretical benefits of previous utilitarian configuration procedures when applied to finite parameter spaces but is significantly faster, both provably and experimentally.

1 INTRODUCTION

Algorithm configuration is a general-purpose technique for optimizing the performance of algorithms by automatically searching over the space of their input parameters. Given a set of possible parameter assignments (i.e., a set of “configurations”) our objective is to find one that makes the algorithm perform well on average, over a set of inputs of interest. Recently, attention has begun to shift from procedures that define performance in terms of average runtime toward procedures that define performance in terms of more general utility functions over runtime. Graham et al. (2023b) introduced Utilitarian Procrastination (UP), the first non-trivial algorithm configuration procedure that maximizes utility achieved instead of minimizing runtime. Unlike more naive approaches, UP can adapt to the hardness of the problem it faces, requiring provably less time to guarantee optimality when many of the configurations it considers are suboptimal. However, UP requires that this set of candidate configurations be discrete and relatively small. This is a major limitation; e.g., many algorithms have at least one continuous parameter. UP cannot effectively search over continuous spaces.

In this paper we present COUP (Continuous, Optimistic Utilitarian Procrastination), which is designed to handle infinite (e.g., continuous) configuration spaces. Like UP, COUP is utilitarian and takes inspiration from the world of multi-armed bandits. However, COUP is based on a different bandit procedure that is guided by the principle of “optimism in the face of uncertainty” and is more readily adapted to handle large configuration spaces, including those with continuous parameters. COUP also offers significantly better performance than UP on finite sets of configurations because COUP does not try to explicitly rule bad configurations out, but instead simply ignores them once they no longer appear as promising as other alternative configurations.

Section 2 discusses relevant background material. Section 3 introduces the finite-configuration-space version of COUP and shows that this version finds an approximately optimal configuration in strictly less total time than UP. We simply call this version OUP since it, like UP, does not actually search over a continuous parameter space. Section 4 introduces COUP, leveraging our understanding of

the finite case into a general-purpose procedure. We discuss empirical evaluations in Section 5 and conclude with Section 6.

2 BACKGROUND

Given a parameterizable algorithm and a set of possible parameter assignments, the goal of *algorithm configuration* is to find a parameter assignment that makes the algorithm perform well across a set of given inputs. If we also require that a procedure gives a *guarantee* about the near-optimality of its returned configuration, algorithm configuration can be seen as a multi-armed bandit problem (see e.g., Lattimore & Szepesvári (2020)) called *approximate best-arm identification*: finding an arm whose mean reward is close to the reward of the best arm available. In this case, it is not the regret of the procedure that concerns us, but the number of samples needed to guarantee the near-optimality of the returned arm. Even-Dar et al. (2002) observe that if we want to find an arm whose value is within ϵ of the best with good probability, we could just take $\tilde{O}(\epsilon^{-2})$ samples of each arm and apply Hoeffding’s inequality (Hoeffding, 1963) to obtain upper and lower confidence bounds on the mean reward of each arm that are sufficient to identify an approximately optimal one (i.e., one within an ϵ additive factor of optimal). However, this will be more samples than necessary in all but the worst case; any arm that is significantly worse than the best arm can be ruled out using fewer samples than are required by naively applying Hoeffding’s inequality. The Successive Elimination (SE) procedure (Even-Dar et al., 2002) is able to identify an approximately optimal arm with only $\tilde{O}(\min\{\Delta_i^{-2}, \epsilon^{-2}\})$ samples of each arm, where Δ_i is the suboptimality gap for arm i . This can be a significant reduction when many arms are far from optimal. SE achieves this improved, *adaptive* bound by being *anytime*: it does not take an ϵ as input. It starts by taking only a few samples of each arm, guaranteeing a relatively large ϵ , and then improves this over time. This allows it to stop sampling bad arms and rule them out before it is finished. Up to logarithmic factors, taking $m = \min\{\Delta_i^{-2}, \epsilon^{-2}\}$ samples of arm i is not only sufficient, but also necessary for general inputs (Mannor & Tsitsiklis, 2004).

SE is a simple and intuitive procedure with approximately optimal sample complexity, but it can tend to over-investigate bad arms, sampling them repeatedly until it can expressly rule them out. The Upper Confidence Bound (UCB) procedure (Lai, 1987) is an alternative to SE that focuses on arms that look promising and/or about which little is known. By repeatedly pulling the arm with largest upper confidence bound, UCB attempts to maximize the information it gains with each pull. For more discussion on the difference between SE and UCB and on their relative performance see the survey by Jamieson & Nowak (2014).

There are two key differences between algorithm configuration and most other bandit problems. First, each arm pull imposes a stochastic cost arising from the amount of time it takes for the algorithm run to complete. Second, we have the option of bounding these costs by “capping” algorithm runs at any fixed time of our choice, but when we do so we obtain only censored information about the arm pull (we learn that the run would have taken more than the captime but not how much more). Overall, it is not the total number of samples that we want to control, but the total amount of time spent searching.

Most existing algorithm configuration procedures that offer theoretical guarantees aim to minimize average runtime. These can broadly be seen as implementing either SE or UCB. These procedures either run all existing configurations simultaneously (i.e., round-robin or in parallel), ruling out the ones that can be proved to be suboptimal after each step (Weisz et al., 2018; 2019; 2020; Brandt et al., 2023), or they iteratively select a particular configuration according to a confidence-bound index, and run only this configuration (Kleinberg et al., 2017; 2019). In either case, these procedures are significantly complicated by the fact that runtimes are potentially unbounded, and even the notion of optimality that they target requires adjustment to compensate for this. Other theoretically-motivated approaches offer performance guarantees based on measures of complexity and guarantee notions of PAC optimality akin to those we present here (Gupta & Roughgarden, 2017; Balcan et al., 2017; 2021). The focus there is not on the time consumed by the configuration process, but on the number of samples required to find a sufficiently good configuration.

Utilitarian algorithm configuration procedures do not try to minimize expected runtime, but instead try to maximize expected utility, as measured by a user-specified utility function that captures the value associated with solving an input instance as a (weakly decreasing, but not necessarily linear) function of the amount of time spent solving it. For example, if we face daily scheduling problems, then getting a solution after waiting 200 days is not 100 times worse than getting it after two days; both are too

long to wait. If one algorithm solves most inputs very quickly but occasionally takes an astronomical amount of time, while another algorithm takes a moderate amount of time on every input, determining which algorithm is “better” will likewise depend on our particular setting: How much time do we have? How much better is it to get an answer quickly? In Graham et al. (2023a) we generalized such intuitions, offering a detailed argument from first principles that utility maximization is a more appropriate objective (e.g., for algorithm configuration) than runtime minimization and showing that such utility functions are guaranteed to exist if our preferences for algorithm runtime distributions follow a certain reasonable structure based on the axioms of Von Neumann & Morgenstern (1947).

UP (Graham et al., 2023b) is a utilitarian configuration procedure that can be seen as SE coupled with a doubling mechanism for discovering sufficiently large captimes. UP has several desirable qualities that we would like to incorporate into our own configuration procedure. First, UP maximizes utility instead of minimizing runtime. Second, unlike many runtime-optimizing procedures UP can make simple theoretical guarantees about the near-optimality of the configuration it returns. Third, UP is *anytime*, meaning that it provides a better guarantee the longer it is run, and so requires minimal initial input from the user, who can simply observe the procedure’s execution and terminate it when they are satisfied with the guarantee being made. Finally, UP’s guaranteed performance is *adaptive* to the inputs it receives; it provably performs better on non-worst-case inputs (i.e., the associated theoretical guarantees are *input-dependent*).

However, UP can only be applied to a finite (and relatively small) set of configurations since it cycles through these one by one and runs each of them in turn. This is a prohibitive limitation we would like to overcome. A natural idea would be to sample a finite set from the infinite configuration space and simply run UP on this set. The optimality guarantee offered by UP would then be made with respect to the best configuration in this subset, and if enough configurations are sampled, we can guarantee that our subset will be likely to contain at least one of the top, say, γ -fraction of all configurations. This strategy will work, but since the choice of γ potentially limits the quality of the configuration we are likely to find, we would like to make our procedure anytime in γ as well; we would like our procedure to start with a relatively small set of configurations (i.e., a relatively large γ) and grow this set as time goes on, shrinking the γ it can guarantee. Of course we still wish to be anytime with respect to the accuracy parameter ϵ , and here we run into a problem with UP. Since UP loops over all remaining configurations, at any point during its execution it will have run all of these configurations on effectively the same number of instances (at most the difference is one). What should we do when we want to expand the set of configurations we are considering? Any newly added configuration will have a large upper confidence bound, since we will know little or nothing about it. This will severely limit the anytime ϵ we can guarantee because it remains possible that the new configuration is much better than all the existing ones. So new configurations would need to be “caught up” when they are added. How should we do this?

One extreme strategy is to completely “catch up” new configurations by running them on all existing instances. This is unnecessarily costly because it fails to leverage what has already been learned about promising configurations. At the other extreme, we could simply refrain from “catching up” new configurations at all, just adding them to the existing pool, forever estimating them with fewer samples and thus lower accuracy. This would effectively limit the ϵ we could guarantee to the lower accuracy of these newly introduced configurations. Is there a way to balance these two extremes and integrate new configurations into the estimation process smoothly? This leads us to the UCB procedure. In terms of the runtime cost we incur, we would prefer to take the second extreme approach above and not do anything special with new configurations. However, this would lead to configurations with relatively large upper confidence bounds, which would limit the optimality guarantee we could make. If we choose instead to focus on configurations that have large upper bounds, as UCB does, we will explore new configurations just up to the point where something else looks better, and then stop.

Thus, whereas UP was built around the SE procedure, COUP is built around the UCB procedure. Like UP, COUP is utilitarian, adaptive, and anytime with respect to the optimality parameter ϵ . At the same time, COUP can also be applied to infinitely-large sets of configurations, and is also anytime with respect to γ , the fraction of unexplored configurations. COUP initially makes guarantees for a small set of configurations, then expands this set while simultaneously improving the optimality guarantee that it can make with respect to this (growing) set. What’s more, COUP also beats UP at its own game: COUP is significantly faster than UP when the input actually is a fixed, finite set of configurations. Because it is built on the SE algorithm, UP runs all suboptimal configurations until they can explicitly be eliminated. For some sets of inputs this can be necessary but in many cases it is

more efficient to rule out poor-performing configurations by learning more about (i.e., by narrowing the confidence bounds of) other, better-performing configurations. COUP provably stops running sub-optimal configurations well before they would be eliminated by UP. Additionally, we re-evaluate a key component of UP—the condition it checks before doubling a configuration’s capttime—and propose an improvement. UP attempts to balance the estimation error due to sampling with the error due to capping. We argue that this is better achieved with a different doubling condition that emerges from a more careful analysis of the confidence bounds.

3 THE CASE OF FEW CONFIGURATIONS: OUP

We begin by presenting and analyzing the finite-configuration-space version of COUP, which we call OUP (Algorithm 1) to indicate that it does not search over continuous spaces like COUP does. OUP can most easily be understood as a UCB procedure for finding configurations with good *capped* mean utility, along with a periodic doubling scheme that allows larger and larger captimes to be considered. Like UP, it is adaptive and anytime in the optimality parameter ϵ .

We are given a set of configurations, indexed by i , and a set of instances, indexed by j . We will use t_{ij} to denote the true uncapped runtime of configuration i on input j , and $t_{ij}(\kappa) := \min(t_{ij}, \kappa)$ to be the κ -capped runtime of i on j . When we do a run, we observe $t_{ij}(\kappa)$ rather than t_{ij} ; these coincide for any run that completes. The instance distribution \mathcal{D}_J , along with any randomness of the algorithm or execution environment will together induce a runtime distribution for each configuration i . We will use \mathcal{D}_i to denote this runtime distribution, and F_i to denote its CDF. For each configuration i , the true uncapped expected utility is $U_i := \mathbb{E}_{t \sim \mathcal{D}_i} [u(t)]$. The cumulative distribution function is $F_i(\kappa) := \Pr_{t \sim \mathcal{D}_i} (t \leq \kappa)$. The capped expected utility is $U_i(\kappa) := \mathbb{E}_{t \sim \mathcal{D}_i} [u(\min(t, \kappa))]$. Finally, the optimality gaps $\Delta_i := \max_{i'} U_{i'} - U_i$ partially determine the hardness of the problem; if a configuration i has large optimality gap Δ_i , then it will be easier to rule out. The notion of optimality that we target is given by the following definition.

Definition 1 (ϵ -optimal). *A configuration i is ϵ -optimal if $\Delta_i \leq \epsilon$.*

The error associated with m runtime samples at capttime κ is $\alpha(m, \kappa) := \sqrt{\frac{\ln(11nm^2(\log \kappa + 1)^2 / \delta)}{2m}}$, chosen to satisfy Hoeffding’s inequality and the necessary union bounds. An execution of OUP will be called *clean* if at all times during its execution we have $|\hat{F}_i - F_i(\kappa_i)| \leq \alpha(m_i, \kappa_i)$ and $|\hat{U}_i - U_i(\kappa_i)| \leq (1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i)$ for all configurations i . So during a clean execution, empirical capped average utilities are close to true capped average utilities and empirical CDF values are close to true CDF values at the capttime κ_i .

Lemma 1. *An execution of OUP is clean with probability at least $1 - \delta$.*

All proofs are deferred to the appendices. The idea behind Lemma 1 is that the bounds that define a clean run hold initially, they do not change in any round where i is not selected, and by Hoeffding’s inequality they hold with high probability in any round where i is selected. The next lemma shows that during a clean execution, the confidence bounds will be valid, and also not too far apart.

Lemma 2. *For all i at all points during a clean execution of OUP we have $LCB_i \leq U_i \leq UCB_i$ and $UCB_i - LCB_i \leq 2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i))$.*

The upper and lower bounds are defined in Algorithm 1. Each iteration of OUP’s main loop corresponds to an iteration of the UCB procedure, with an additional check to see if the capttime needs to be doubled. In each iteration the configuration with largest upper bound is selected (Line 7). The doubling condition is then checked (Line 9) and, if necessary, instances that had previously capped are rerun at the newly doubled capttime (Line 11). A new runtime sample is drawn, and the sample mean and confidence bounds are recomputed. Finally, we attempt to eliminate provably suboptimal configurations (Lines 20-24). This loop repeats until terminated by the user, at which point it returns the configuration with largest lower confidence bound, which is the configuration we can prove the smallest suboptimality for. We can now state the main theorem of this section.

Theorem 1. *With probability at least $1 - \delta$, OUP eventually returns an optimal configuration and it returns an ϵ -optimal configuration if it is run until $2\alpha_{i_{opt}} + u(\kappa_{i_{opt}})(1 - F_{i_{opt}}(\kappa_{i_{opt}})) \leq \epsilon$. Furthermore, for any suboptimal configuration i , if m_i and κ_i ever become large enough that*

$$2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) < \Delta_i \quad (1)$$

Algorithm 1 OUP

```

1: Input: configurations  $i = 1, \dots, n$ ; instances  $j = 1, 2, \dots$ ; utility function  $u$ ; failure probability  $\delta$ .
2:  $I \leftarrow \{1, \dots, n\}$  ▷ candidate configurations
3: for  $i \in I$  do
4:    $LCB_i \leftarrow 0$ ;  $UCB_i \leftarrow 1$ ;  $\hat{U}_i \leftarrow 0$ ;  $\hat{F}_i \leftarrow 0$ ;  $m_i \leftarrow 0$ ;  $\kappa_i \leftarrow 1$  ▷ initializations
5: end for
6: for  $r = 1, 2, 3, \dots$  do
7:    $i \leftarrow \arg \max_{i' \in I} UCB_{i'}$ 
8:    $m_i \leftarrow m_i + 1$ 
9:   if  $2\alpha(m_i, \kappa_i) \leq u(\kappa_i)(1 - \hat{F}_i)$  then ▷ captime doubling condition
10:     $\kappa_i \leftarrow 2\kappa_i$ 
11:     $t_{i1}(\kappa_i), \dots, t_{im_i}(\kappa_i) \leftarrow$  runtime of configuration  $i$  on instances  $1, \dots, m_i$  with timeout  $\kappa_i$ 
12:   else
13:     $t_{im_i}(\kappa_i) \leftarrow$  runtime of configuration  $i$  on instance  $m_i$  with timeout  $\kappa_i$ 
14:   end if
15:    $\hat{F}_i \leftarrow \frac{|\{j \in [m_i] : t_{ij}(\kappa_i) < \kappa_i\}|}{m_i}$  ▷ fraction of runs that completed
16:    $\hat{U}_i \leftarrow \frac{1}{m_i} \sum_{j=1}^{m_i} u(t_{ij}(\kappa_i))$  ▷ empirical average utility
17:    $UCB_i \leftarrow \hat{U}_i + (1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i)$ 
18:    $LCB_i \leftarrow \hat{U}_i - \alpha(m_i, \kappa_i) - u(\kappa_i)(1 - \hat{F}_i)$ 
19:    $i^* \leftarrow \arg \max_{i' \in I} LCB_{i'}$ 
20:   for  $i' \in I$  do
21:     if  $UCB_{i'} < LCB_{i^*}$  then ▷ if  $i'$  is suboptimal
22:        $I \leftarrow I \setminus \{i'\}$  ▷ remove  $i'$  from consideration
23:     end if
24:   end for
25:   if execution is terminated or  $|I| = 1$  then
26:     return  $i^*$ 
27:   end if
28: end for

```

then i will never be run again, and i will be outright eliminated once $m_i, \kappa_i, m_{i^{opt}}$ and $\kappa_{i^{opt}}$ are large enough that

$$2\alpha(m_i, \kappa_i) + 2\alpha(m_{i^{opt}}, \kappa_{i^{opt}}) + u(\kappa_i)(1 - F_i(\kappa_i)) + u(\kappa_{i^{opt}})(1 - F_{i^{opt}}(\kappa_{i^{opt}})) < \Delta_i. \quad (2)$$

If OUP keeps selecting and running i , then the term $2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i))$ will continue to shrink. So we will eventually have $2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) < \Delta_i$ as long as we keep running i . In either case, we will eventually stop running any suboptimal i . This puts a hard limit on the amount of time we will ever dedicate to any suboptimal configuration, and this limit is better than the corresponding guarantee made by UP: compare our Theorem 1 to Theorem 4 in Graham et al. (2023b). UP guarantees that i will be eliminated once the bound in Equation (2) is satisfied. The bound in Equation (1) guaranteed by OUP is a necessary condition for this. OUP does not keep running suboptimal configurations until they can be eliminated, as UP does. Instead, it simply stops running them, while continuing to tighten the bounds of better configurations instead.

Algorithm 1 is stated using the old doubling condition introduced in UP. We propose an alternative based on examining the confidence width $UCB_i - LCB_i = 2(1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i) + u(\kappa_i)(1 - \hat{F}_i + \alpha(m_i, \kappa_i))$. The first term is the error from runs below κ_i and the second term is the error from runs above κ_i . Our new doubling condition attempts to balance these two terms. Due to space constraints we discuss this more in Appendix B, but ultimately Line 9 of Algorithm 1 becomes

9': **if** $2(1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i) \leq u(\kappa_i)(1 - \hat{F}_i + \alpha(m_i, \kappa_i))$ **then** ▷ captime doubling condition

The change to the continuous-space version COUP is analogous (Line 18 of Algorithm 2). Implementing this refined doubling condition improves the performance of both UP and (C)OUP.

4 THE CASE OF MANY CONFIGURATIONS: COUP

We now assume we have a very large, possibly uncountably infinite, set of configurations \mathcal{A} , along with an associated distribution $\mathcal{D}_{\mathcal{A}}$ over configurations $a \in \mathcal{A}$ from which we sample. It is no longer feasible to target the configuration with maximum utility $OPT = \max_{a \in \mathcal{A}} U_a$. It may be arbitrarily unlikely that we ever sample this configuration and, in fact, such a maximum may not even exist. So we must relax our objective. For any $\gamma \in (0, 1)$, let $OPT^\gamma = \sup \{ \mu : \Pr_{a \sim \mathcal{D}_{\mathcal{A}}} [U_a \leq \mu] \leq 1 - \gamma \}$. The quantity OPT^γ marks the top $\lfloor 1/\gamma \rfloor$ -quantile of expected utilities. That is, OPT^γ is the utility of the best configuration that remains after we have excluded the top γ -fraction of configurations. The notion of optimality that we will target is given by the following definition.

Definition 2 ((ϵ, γ) -optimal). A configuration $a \in \mathcal{A}$ is (ϵ, γ) -optimal if $U_a \geq OPT^\gamma - \epsilon$.

We present the general version of COUP in Algorithm 2. The procedure runs in successive phases. Each phase is an execution of the few-configuration version, OUP, with the bounds being chosen carefully so that runs can be shared across phases. Each phase p of COUP is characterized by two parameters: $\gamma_p \in (0, 1)$ determines the tail quantile being targeted in this phase, and $\epsilon_p \in (0, 1)$ determines the level of optimality being targeted. At the start of each phase p , COUP samples enough new configurations to ensure that it has a total of $n_p = \frac{1}{\gamma_p} \ln \frac{\pi^2 p^2}{3\delta}$ configurations, where δ is the failure probability and π is the mathematical constant (Lines 5 and 6). COUP then essentially runs the finite-space version OUP on this set of n_p configurations until it can prove ϵ_p -optimality (Lines 15–28). However, COUP “maintains state” between phases so that it can build on what it has already learned. Confidence bounds in phase p incorporate any runs performed in previous phases (Lines 12 and 13), with the error terms being set carefully so that the total failure probability is controlled across all phases.

If γ_p is relatively small and ϵ_p is relatively large, then COUP will sample a large number of configurations in phase p , but not work too hard to prove that any of them is approximately optimal. On the other hand, if γ_p is relatively large and ϵ_p is relatively small, then COUP will sample only a few configurations and work very hard to prove one of them is nearly optimal (with respect to this set). Since we are interested in procedures that are anytime in both ϵ and γ , we will be interested in sequences of parameters for which $\epsilon_p \rightarrow 0$ and $\gamma_p \rightarrow 0$ as $p \rightarrow \infty$. We can think of ϵ_p as controlling the rate at which we explore new instances, and γ_p as controlling the rate at which we explore new configurations. They are dissimilar in what they quantify though: ϵ_p is in “units” of utility, while γ_p is in “units” of probability.

COUP never eliminates configurations because it needs them in later phases to make statistical guarantees about relative optimality, but we know it will stop running them well before we will be able to consider them for elimination anyway. Define $\alpha_p(m, \kappa) := \sqrt{\frac{\ln(36p^2 n_p m^2 (\log \kappa + 1)^2 / \delta)}{2m}}$ to be the confidence width, chosen to satisfy Hoeffding’s inequality and the necessary union bounds. The p -th phase of an execution of COUP will be called *clean* if $|\hat{F}_i - F_i(\kappa_i)| \leq \alpha_p(m_i, \kappa_i)$ and $|\hat{U}_i - U_i(\kappa_i)| \leq (1 - u(\kappa_i)) \cdot \alpha_p(m_i, \kappa_i)$ for all m_i and κ_i for all configurations $i \in \mathcal{A}_p$.

Lemma 3. An execution of COUP is clean for all phases with probability at least $1 - \frac{\delta}{2}$.

Lemma 4. During a clean execution of COUP we have $LCB_i \leq U_i \leq UCB_i$ and $UCB_i - LCB_i \leq 2\alpha_p(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i))$ for all configurations $i \in \mathcal{A}_p$ in all phases p .

The p -th phase of an execution of COUP will be called *characteristic* if there exists some $i_p \in \mathcal{A}_p$ with $U_{i_p} \geq OPT^{\gamma_p}$.

Lemma 5. An execution of COUP is characteristic for all phases with probability at least $1 - \frac{\delta}{2}$.

Finally, we can state our main theorem about COUP’s performance.

Theorem 2. If COUP is run with parameters $(\epsilon_1, \epsilon_2, \dots)$ and $(\gamma_1, \gamma_2, \dots)$ then with probability at least $1 - \delta$ it returns an (ϵ_p, γ_p) -optimal configuration at the end of every phase $p = 1, 2, \dots$

5 EMPIRICAL EVALUATION

We first compare the performance of OUP with the performance of UP (Section 5.1), and with a naive procedure also considered in our previous work (Graham et al., 2023b) that simply takes

Algorithm 2 COUP

```

1: Input: distribution over configurations  $\mathcal{D}_{\mathcal{A}}$ ; instances  $j = 1, 2, \dots$ ; utility function  $u$ ; failure
   parameter  $\delta$ ; phase parameters  $\{\epsilon_p, \gamma_p\}_{p=1,2,3,\dots}$ 
2:  $n_0 \leftarrow 0$ 
3:  $\mathcal{A}_0 \leftarrow \emptyset$ 
4: for  $p = 1, 2, 3, \dots$  until terminated do
5:    $n_p \leftarrow \left\lceil \frac{\ln \frac{\pi^2 p^2}{3\delta}}{\gamma_p} \right\rceil$ 
6:    $N_p \leftarrow \text{sample } n_p - |\mathcal{A}_{p-1}| \text{ new configurations from } \mathcal{D}_{\mathcal{A}}$ 
7:    $\mathcal{A}_p \leftarrow \mathcal{A}_{p-1} \cup N_p$  ▷ all configurations for this phase
8:   for  $i \in N_p$  do ▷ initializations for new configurations
9:      $LCB_i \leftarrow 0, UCB_i \leftarrow 1; \hat{U}_i \leftarrow 0; \hat{F}_i \leftarrow 0; m_i \leftarrow 0; \kappa_i \leftarrow 1$ 
10:  end for
11:  for  $i \in \mathcal{A}_{p-1}$  do ▷ update bounds for existing configurations
12:     $UCB_i \leftarrow \hat{U}_i + (1 - u(\kappa_i)) \cdot \alpha_p(m_i, \kappa_i)$ 
13:     $LCB_i \leftarrow \hat{U}_i - \alpha_p(m_i, \kappa_i) - u(\kappa_i)(1 - \hat{F}_i)$ 
14:  end for
15:  while  $\max_{i \in \mathcal{A}_p} UCB_i - \max_{i \in \mathcal{A}_p} LCB_i \geq \epsilon_p$  do ▷ phase termination condition
16:     $i \leftarrow \arg \max_{i' \in \mathcal{A}_p} UCB_{i'}$ 
17:     $m_i \leftarrow m_i + 1$ 
18:    if  $2\alpha_p(m_i, \kappa_i) \leq u(\kappa_i)(1 - \hat{F}_i)$  then ▷ capttime doubling condition
19:       $\kappa_i \leftarrow 2\kappa_i$ 
20:       $t_{i1}(\kappa_i), \dots, t_{im_i}(\kappa_i) \leftarrow \text{runtime of } i \text{ on instances } 1, \dots, m_i \text{ with timeout } \kappa_i$ 
21:    else
22:       $t_{im_i}(\kappa_i) \leftarrow \text{runtime of } i \text{ on instance } m_i \text{ with timeout } \kappa_i$ 
23:    end if
24:     $\hat{F}_i \leftarrow \frac{|\{j \in [m_i] : t_{ij}(\kappa_i) < \kappa_i\}|}{m_i}$  ▷ fraction of runs that completed
25:     $\hat{U}_i \leftarrow \frac{1}{m_i} \sum_{j=1}^{m_i} u(t_{ij}(\kappa_i))$  ▷ empirical average utility
26:     $UCB_i \leftarrow \hat{U}_i + (1 - u(\kappa_i)) \cdot \alpha_p(m_i, \kappa_i)$ 
27:     $LCB_i \leftarrow \hat{U}_i - \alpha_p(m_i, \kappa_i) - u(\kappa_i)(1 - \hat{F}_i)$ 
28:  end while
29: end for
30: return  $\arg \max_{i \in \mathcal{A}_p} LCB_i$ 

```

enough runtime samples to satisfy Hoeffding’s inequality for each configuration, and takes them at a high enough capttime that it can still make its guarantee whether they complete or not. We then explore the effects on performance when we expand our search to an infinitely-large set of configurations (Section 5.2). We show that the search of the configuration space does not cost COUP too much in terms of runtime, and we demonstrate COUP’s flexibility in the way it performs its search. We perform our experiments on three existing datasets from the algorithm configuration literature. The `minisat` dataset represents the measured execution times of the SAT solver by that name on a generated set of instances (see Weisz et al. (2018) for details). The `cplex_rcw` and `cplex_region` datasets represent the predicted solve times of the CPLEX integer program solver on wildlife conservation and combinatorial auction problems, respectively (see Weisz et al. (2020) for details). In Section 5.3 we compare COUP and OUP to a variety of other procedures that do not make utilitarian optimality guarantees. We focus on the two utility functions considered in our previous work (Graham et al., 2023b). The “log-Laplace” utility function is defined as $u_{LL}(t) = 1 - \frac{1}{2} \left(\frac{t}{\kappa_0}\right)^\alpha$ if $t < \kappa_0$ and $u_{LL}(t) = \frac{1}{2} \left(\frac{\kappa_0}{t}\right)^\alpha$ otherwise. The “uniform” utility function is defined as $u_{unif}(t) = 1 - \frac{t}{\kappa_0}$ if $t < \kappa_0$ and $u_{unif}(t) = 0$ otherwise. The parameter κ_0 is set to 60 in both cases, and the parameter α is set to 1. To facilitate the comparison with UP we use the original doubling condition when recreating our previous experiments (Graham et al., 2023b, Figures 1 and 2). For other experiments we use the improved doubling condition. We perform all experiments in an environment where restarting runs is not possible. Code to reproduce all experiments can be found at <https://github.com/drgrhm/coup>.

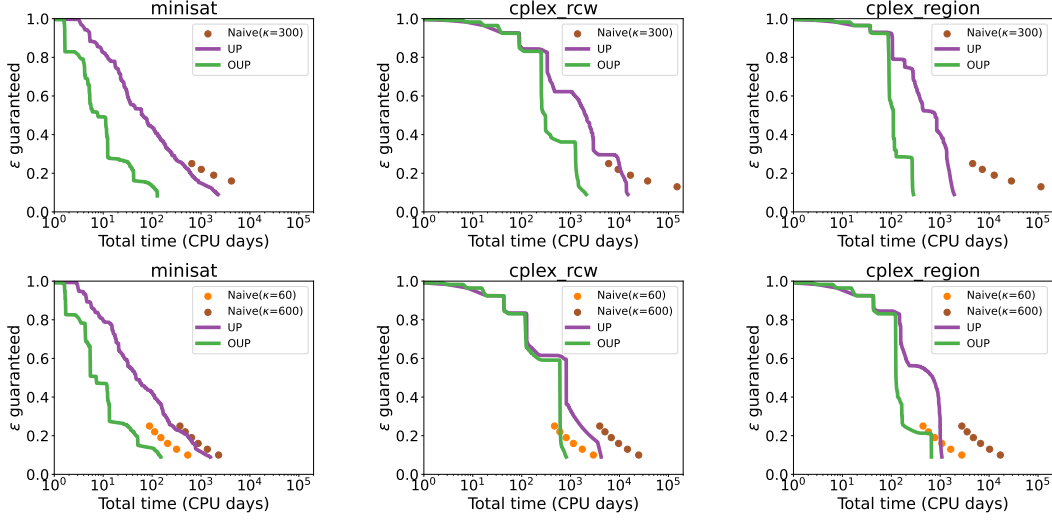


Figure 1: The ϵ guaranteed by each procedure as a function of total runtime using the log-Laplace utility function (top row) and uniform utility function (bottom row). OUP consistently outperforms both UP and the naive procedure for reasonable values of ϵ , often by an order of magnitude and even when the the parameter of the naive procedure has been well-chosen.




Figure 2: Total time spent by OUP on each configuration (log-Laplace utility function). Configurations are sorted according to average utility. OUP spends less time on all but the very best configurations.

5.1 THE CASE OF FEW CONFIGURATIONS

We first compare the performance of OUP with that of UP by recreating our previously published experiments (Graham et al., 2023b). Figure 1 shows that OUP consistently proves a meaningful ϵ in much less time than UP. OUP is faster by an order of magnitude or more (compare the top row of Figure 1 to Figure 2 in Graham et al. (2023b)). Furthermore, a very simple naive approach based on Hoeffding’s inequality can be faster than UP in some scenarios (i.e. for relatively small values of ϵ) when the capttime parameter is chosen appropriately for the given utility function. However, OUP is consistently faster than both UP and this naive approach, regardless of the chosen capttime parameter. In Section 3 we argued that OUP will spend less time running bad configurations than UP because it can stop devoting attention to them and focus on more promising ones instead. UP, on the other hand, continues to run suboptimal configurations until they are eliminated. Figure 2 shows that OUP does indeed spend considerably less time on suboptimal configurations.

5.2 THE CASE OF MANY CONFIGURATIONS

We now compare the performance of COUP with that of OUP to show that COUP’s exploration of the configuration space does not cost too much. The results for the log-Laplace utility function are shown in Figure 3. We set $\delta = 0.01$ throughout. For each phase p we set $\epsilon_p = e^{-p/6}$ and $\gamma_p = e^{-p/3}$. This schedule allows COUP to sample a large fraction of the configurations and the instances in the datasets we consider. When it is finished each phase, it can make an ϵ_p -optimality guarantee with respect to the set \mathcal{A}_p of n_p configurations. We compare the total configuration time of COUP at

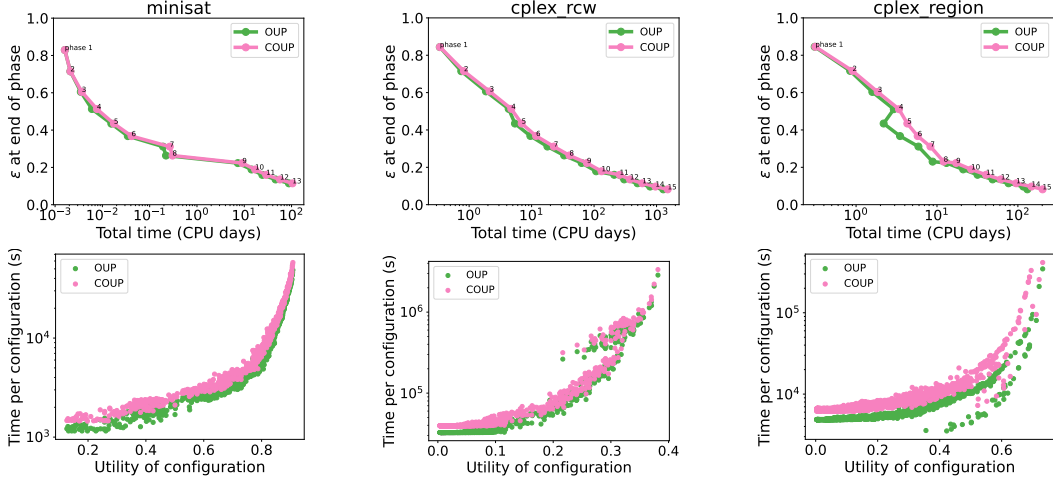


Figure 3: Performance of COUP compared to OUP, using the log-Laplace utility function. COUP is anytime, but takes only a small factor more time than OUP. The bottom row shows the final times. OUP backtracks on the `cplex_region` dataset because COUP sampled a relatively good configuration in phase 5, which OUP is able to capitalize on, proving optimality in less total time than in phase 4.




Figure 4: Different schedules for exploring new configurations (refining γ) vs. exploring existing configurations (refining ϵ) when running COUP. Results are for the log-Laplace utility function and the `cplex_rcw` dataset.

the end of each phase to the time OUP takes to prove ϵ_p -optimality when applied directly to this set of n_p configurations. Figure 3 shows that the total configuration time required by COUP is not much different from the time required by OUP, indicating that COUP was efficient at exploring the configuration space.

COUP is extremely flexible in the way it explores new configurations relative to exploring new instances. The parameter sequences $\gamma_1, \gamma_2, \dots$ and $\epsilon_1, \epsilon_2, \dots$ control the relative rates at which this exploration takes place. We demonstrate this flexibility in Figure 4, presenting various different exploration strategies. The “ γ focus” strategy sets $\epsilon_p = e^{-p/30}$ and $\gamma_p = e^{-p/3}$, and so works to explore many configurations, making relatively weak guarantees about their optimality. The “ ϵ focus” strategy sets $\epsilon_p = e^{-p/3}$ and $\gamma_p = e^{-p/30}$, and so explores relatively few configurations, but makes strong optimality guarantees with respect to these. The “balanced” strategy sets $\epsilon_p = \gamma_p = e^{-p/5}$. In general, it is not wise to focus on ϵ too intently until later in the search process. Once we have found a good configuration (i.e., one with a relatively large upper confidence bound), new configurations can be ruled out relatively cheaply, because they quickly look worse than the good one. So it is helpful to find a good configuration quickly, and then later work on proving that it is good. The “ γ then ϵ ” strategy achieves this by setting $\epsilon_p = e^{-p^3/300}$ and $\gamma_p = e^{-p^2/30}$.

5.3 COMPARISON TO PROCEDURES WITH DIFFERENT GUARANTEES

Finally, while we see the main contribution of COUP as being its ability to search infinite parameter spaces and make utilitarian near-optimality guarantees with respect to them, which other configuration procedures have not been able to do, some other procedures do make different but related guarantees. Comparing the performance of these various methods is not always easy since they

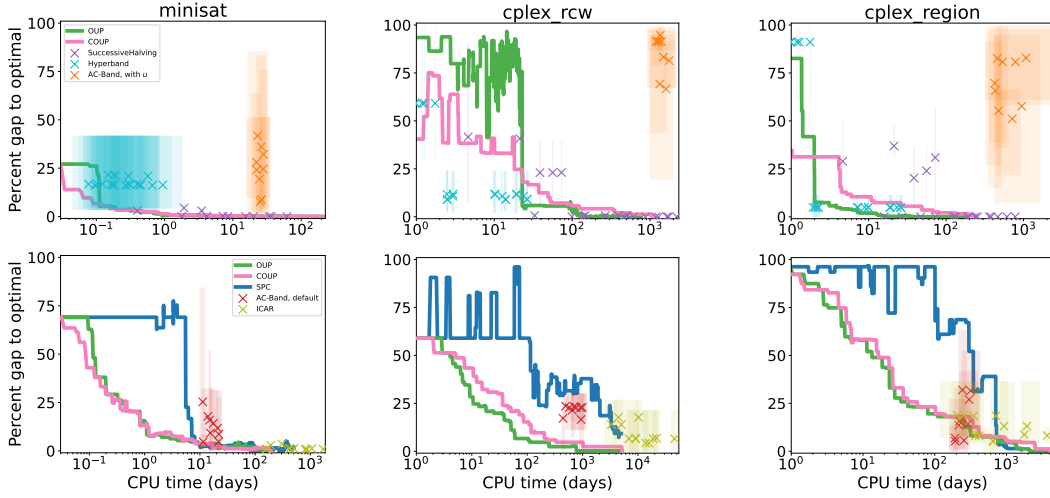


Figure 5: Average configurator performance as measured by the percentage gap to best configuration in the dataset. Top row shows procedures optimizing the log-Laplace utility function. Bottom row shows procedures optimizing an expected runtime. Error regions show maximum and minimum observations.

optimize different metrics and offer different guarantees. We compare COUP and OUP to Successive Halving (Jamieson & Talwalkar, 2016), Hyperband (Li et al., 2018), AC-Band (Brandt et al., 2023), ImpatientCapsAndRuns (ICAR) (Weisz et al., 2020) and Structured Procrastination with Confidence (SPC) (Kleinberg et al., 2019). We describe each of these procedures and their parameter settings in more detail in the appendices. Figure 5 shows the percentage gap to best configuration for each procedure at various amounts of runtime. The top row shows utilitarian procedures optimized for the log-Laplace utility function, while the bottom row shows procedures that optimize runtime. In the latter case, we optimized OUP and COUP on a uniform utility function, which is equivalent to optimizing κ_0 -capped runtime. For the `minisat` dataset we set $\kappa_0 = 100$ and for the CPLEX datasets we set $\kappa_0 = 3000$. These values are similar to the per-configuration captimes used by ICAR and SPC, which both optimize an expected quantile-capped runtime. Figure 5 shows the results for a quantile parameter of $\delta = .1$. The results for $\delta = .2$ are essentially the same (see Appendix C). To avoid an unfair comparison we report the performance of OUP and COUP with respect to this quantile-capped runtime metric. We observed that AC-Band performed poorly with log-Laplace utility functions, so we additionally report its performance gap with respect to its default metric. We plot the averages and error regions over five seeds. OUP, COUP and SPC are anytime so they make recommendations throughout their execution, while the other procedures make point recommendations. Overall, COUP and OUP performed better than many runs of the other procedures and eventually narrowed in on the optimal configuration for every seed. Hyperband also worked well by this percentage-gap metric in some settings on some seeds. For the RCW dataset it found a good configuration more quickly than COUP but still left a gap to the best, whereas COUP and OUP continually improved. COUP and OUP avoided the “luck of the draw” that other non-anytime methods faced in the choice of their budget parameters; our investigation of varying parameterizations for these procedures shows that these hard-to-set parameters substantially impacted performance.

6 CONCLUSION

We have presented COUP, an improved method for utilitarian algorithm configuration. COUP shares the positive qualities of its predecessors, but is truly general purpose, in that it can search over infinite configuration spaces effectively. COUP is also shown to be superior when applied to finite configuration spaces. For the sake of continuity with previous work we have made comparisons on the datasets and utility functions that have been used before. However, this comparison is somewhat limited; a more comprehensive investigation would be a valuable direction for future work. Moreover, while we are convinced that optimizing utility functions is better than optimizing runtime, we recognize that coming up with the right utility function is not always easy for end users; establishing a tool or guide for doing this is another important future direction.

ACKNOWLEDGMENTS

This work was funded by an NSERC Discovery Grant, a CIFAR Canada AI Research Chair (Alberta Machine Intelligence Institute), and computational resources provided both by UBC Advanced Research Computing and a Digital Research Alliance of Canada RAC Allocation.

REFERENCES

- Maria-Florina Balcan, Vaishnavh Nagarajan, Ellen Vitercik, and Colin White. Learning-theoretic foundations of algorithm configuration for combinatorial partitioning problems. In *Conference on Learning Theory*, pp. 213–274, 2017.
- Maria-Florina Balcan, Dan DeBlasio, Travis Dick, Carl Kingsford, Tuomas Sandholm, and Ellen Vitercik. How much data is sufficient to learn high-performing algorithms? generalization guarantees for data-driven algorithm design. In *Symposium on Theory of Computing*, pp. 919–932, 2021.
- Jasmin Brandt, Elias Schede, Björn Haddendorst, Viktor Bengs, Eyke Hüllermeier, and Kevin Tierney. AC-Band: a combinatorial bandit-based approach to algorithm configuration. In *AAAI Conference on Artificial Intelligence*, pp. 12355–12363, 2023.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. PAC bounds for multi-armed bandit and markov decision processes. In *Conference on Computational Learning Theory*, pp. 255–270. Springer, 2002.
- Devon R Graham, Kevin Leyton-Brown, and Tim Roughgarden. Formalizing preferences over runtime distributions. In *International Conference on Machine Learning*, pp. 11659–11682, 2023a.
- Devon R Graham, Kevin Leyton-Brown, and Tim Roughgarden. Utilitarian algorithm configuration. In *Neural Information Processing Systems*, pp. 11659–11682, 2023b.
- Rishi Gupta and Tim Roughgarden. A PAC approach to application-specific algorithm selection. *SIAM Journal on Computing*, 46(3):992–1017, 2017.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. ISSN 01621459.
- Kevin Jamieson and Robert Nowak. Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *Conference on Information Sciences and Systems*, pp. 1–6. IEEE, 2014.
- Kevin Jamieson and Ameet Talwalkar. Non-stochastic best arm identification and hyperparameter optimization. In *Artificial intelligence and statistics*, pp. 240–248. PMLR, 2016.
- Robert Kleinberg, Kevin Leyton-Brown, and Brendan Lucier. Efficiency through procrastination: Approximately optimal algorithm configuration with runtime guarantees. In *International Joint Conference on Artificial Intelligence*, pp. 2023–2031, 2017.
- Robert Kleinberg, Kevin Leyton-Brown, Brendan Lucier, and Devon Graham. Procrastinating with confidence: Near-optimal, anytime, adaptive algorithm configuration. In *Neural Information Processing Systems*, pp. 8881–8891, 2019.
- Tze Leung Lai. Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics*, pp. 1091–1114, 1987.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, 18(185):1–52, 2018.
- Shie Mannor and John N Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun):623–648, 2004.

John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton university press, 2nd edition, 1947.

Gellért Weisz, Andras Gyorgy, and Csaba Szepesvári. LeapsAndBounds: a method for approximately optimal algorithm configuration. In *International Conference on Machine Learning*, pp. 5257–5265, 2018.

Gellért Weisz, Andras Gyorgy, and Csaba Szepesvári. CapsAndRuns: an improved method for approximately optimal algorithm configuration. In *International Conference on Machine Learning*, pp. 6707–6715, 2019.

Gellért Weisz, András György, Wei-I Lin, Devon Graham, Kevin Leyton-Brown, Csaba Szepesvari, and Brendan Lucier. ImpatientCapsAndRuns: approximately optimal algorithm configuration from an infinite pool. In *Neural Information Processing Systems*, pp. 17478–17488, 2020.

A DEFERRED PROOFS

Lemma 1: *An execution of OUP is clean with probability at least $1 - \delta$.*

Proof. For any configuration i and any round r , let $m_i(r)$ and $\kappa_i(r)$ be the number of samples taken by i and the captime they were taken with, respectively, in round r . Let $\widehat{F}_i(m_i, \kappa_i)$ and $\widehat{U}_i(m_i, \kappa_i)$ be OUP’s internal variables \widehat{F}_i and \widehat{U}_i after taking m_i samples at captime κ_i . Define the good events:

$$\begin{aligned}\mathcal{G}_{i,m_i,\kappa_i}^{(F)} &= \{|\widehat{F}_i(m_i, \kappa_i) - F_i(\kappa_i)| \leq \alpha(m_i, \kappa_i)\} \\ \mathcal{G}_{i,m_i,\kappa_i}^{(U)} &= \{|\widehat{U}_i(m_i, \kappa_i) - U_i(\kappa_i)| \leq (1 - u(\kappa_i))\alpha(m_i, \kappa_i)\} \\ \mathcal{G}_{i,m_i,\kappa_i} &= \mathcal{G}_{i,m_i,\kappa_i}^{(F)} \cap \mathcal{G}_{i,m_i,\kappa_i}^{(U)} \\ \mathcal{G}_{i,r} &= \mathcal{G}_{i,m_i(r),\kappa_i(r)} \\ \mathcal{G}_i &= \bigcap_{r=1}^{\infty} \mathcal{G}_{i,r} \\ \mathcal{G} &= \mathcal{G}_1 \cap \dots \cap \mathcal{G}_n.\end{aligned}$$

By definition, an execution is clean if and only if \mathcal{G} holds. Due to the captime doubling condition, and because i might eventually be eliminated (though it might not be), if R_i is the set of rounds in which i was selected, we have

$$\bigcap_{m_i=1}^{\infty} \bigcap_{l_i=1}^{\infty} \mathcal{G}_{i,m_i,2^{l_i-1}} \subseteq \bigcap_{r \in R_i} \mathcal{G}_{i,m_i(r),\kappa_i(r)}. \quad (3)$$

For any round r , let $\rho_i(r) \leq r$ be the last round in which configuration i was selected (i.e., up to and including round r). Since i was not selected in any rounds between $\rho_i(r)$ and r , the values of i ’s internal variables in round r are the same as they were in round $\rho_i(r)$. So the clean bounds hold in round r if and only if they hold in round $\rho_i(r)$, and so they hold in all rounds if and only if they hold in all rounds where i is selected. In other words

$$\bigcap_{r=1}^{\infty} \mathcal{G}_{i,r} = \bigcap_{r \in R_i} \mathcal{G}_{i,r}. \quad (4)$$

We can now bound the probability that an execution is not clean:

$$\begin{aligned}
\Pr(\bar{\mathcal{G}}) &= \Pr(\bar{\mathcal{G}}_1 \cup \dots \cup \bar{\mathcal{G}}_n) && \text{(de Morgan's law)} \\
&\leq \sum_{i=1}^n \Pr(\bar{\mathcal{G}}_i) && \text{(union bound)} \\
&= \sum_{i=1}^n \Pr\left(\overline{\bigcap_{r=1}^{\infty} \mathcal{G}_{i,r}}\right) && \text{(definition)} \\
&= \sum_{i=1}^n \Pr\left(\overline{\bigcap_{r \in R_i} \mathcal{G}_{i,r}}\right) && \text{(Equation (4))} \\
&= \sum_{i=1}^n \Pr\left(\overline{\bigcap_{r \in R_i} \mathcal{G}_{i,m_i(r),\kappa_i(r)}}\right) && \text{(definition)} \\
&\leq \sum_{i=1}^n \Pr\left(\overline{\bigcap_{m_i=1}^{\infty} \bigcap_{l_i=1}^{\infty} \mathcal{G}_{i,m_i,2^{l_i-1}}}\right) && \text{(Equation (3))} \\
&= \sum_{i=1}^n \Pr\left(\overline{\bigcup_{m_i=1}^{\infty} \bigcup_{l_i=1}^{\infty} \mathcal{G}_{i,m_i,2^{l_i-1}}}\right) && \text{(de Morgan's law)} \\
&\leq \sum_{i=1}^n \sum_{m_i=1}^{\infty} \sum_{l_i=1}^{\infty} \Pr(\bar{\mathcal{G}}_{i,m_i,2^{l_i-1}}) && \text{(union bound)} \\
&\leq \sum_{i=1}^n \sum_{m_i=1}^{\infty} \sum_{l_i=1}^{\infty} \left(\Pr(\bar{\mathcal{G}}_{i,m_i,\kappa_i}^{(F)}) + \Pr(\bar{\mathcal{G}}_{i,m_i,\kappa_i}^{(U)}) \right) && \text{(union bound)} \\
&\leq \sum_{i=1}^n \sum_{m_i=1}^{\infty} \sum_{l_i=1}^{\infty} \frac{4\delta}{11nm_i^2l_i^2} && \text{(Hoeffding's inequality)} \\
&\leq \delta && \text{(Basel problem).}
\end{aligned}$$

□

Lemma 2: For all i at all points during a clean execution, we have

$$LCB_i \leq U_i \leq UCB_i,$$

and

$$UCB_i - LCB_i \leq 2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)).$$

Proof. We will use the fact that $U_i(\kappa) - u(\kappa)(1 - F_i(\kappa)) \leq U_i$ for any κ by the law of total expectation, and that $U_i \leq U_i(\kappa)$ by the monotonicity of u . We have

$$\begin{aligned}
LCB_i &= \hat{U}_i - \alpha(m_i, \kappa_i) - u(\kappa_i)(1 - \hat{F}_i) && \text{(definition)} \\
&\leq U_i(\kappa_i) - u(\kappa_i) \cdot \alpha(m_i, \kappa_i) - u(\kappa_i)(1 - \hat{F}_i) && \text{(clean)} \\
&\leq U_i(\kappa_i) - u(\kappa_i)(1 - F_i(\kappa_i)) && \text{(clean)} \\
&\leq U_i && \text{(law of total expectation)} \\
&\leq U_i(\kappa_i) && \text{(monotonic utility)} \\
&\leq \hat{U}_i + (1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i) && \text{(clean)} \\
&= UCB_i && \text{(definition).}
\end{aligned}$$

And

$$\begin{aligned}
UCB_i - LCB_i &= (2 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i) + u(\kappa_i)(1 - \hat{F}_i) && \text{(definition)} \\
&\leq 2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) && \text{(clean).}
\end{aligned}$$

□

Theorem 1: *With probability at least $1 - \delta$, OUP eventually returns an optimal configuration and it returns an ϵ -optimal configuration if it is run until $2\alpha_{i^{opt}} + u(\kappa_{i^{opt}})(1 - F_{i^{opt}}(\kappa_{i^{opt}})) \leq \epsilon$. Furthermore, for any suboptimal configuration i , if m_i and κ_i ever become large enough that*

$$2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) < \Delta_i$$

then i will never be run again, and i will be outright eliminated once $m_i, \kappa_i, m_{i^{opt}}$ and $\kappa_{i^{opt}}$ are large enough that

$$2\alpha(m_i, \kappa_i) + 2\alpha(m_{i^{opt}}, \kappa_{i^{opt}}) + u(\kappa_i)(1 - F_i(\kappa_i)) + u(\kappa_{i^{opt}})(1 - F_{i^{opt}}(\kappa_{i^{opt}})) < \Delta_i.$$

Proof. By Lemma 1, the execution is clean with probability at least $1 - \delta$. Suppose the execution is clean, so the bounds in Lemma 2 hold. For any configuration i in any round r , suppose that m_i and κ_i are large enough that $2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) < \Delta_i$ at the beginning of round r (i.e., before any runs are performed). We have

$$\begin{aligned} UCB_i &\leq LCB_i + 2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) && \text{(Lemma 2)} \\ &\leq U_i + 2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) && \text{(Lemma 2)} \\ &< U_i + \Delta_i && \text{(assumption)} \\ &\leq U_{i^{opt}} && \text{(definition)} \\ &\leq UCB_{i^{opt}} && \text{(Lemma 2)} \end{aligned}$$

so i will not be selected in round r . And since $2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i))$ does not change in any round that i is not selected, i will never be selected again. So if $2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) < \Delta_i$, then i will never be run again.

Now, if OUP keeps selecting and running i , then we will have $2\alpha(m_i, \kappa_i) \rightarrow 0$ and $u(\kappa_i)(1 - F_i(\kappa_i)) \rightarrow 0$, so there will come a time where $2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) < \Delta_i$. In either case there will be some final round r_i after which i is never run again. Let m_i and κ_i be the procedure's internal parameters at the end of this round. Only the optimal configuration(s) will be run after round $\max_{i \text{ suboptimal}} r_i$. Define $\gamma_i = \Delta_i - 2\alpha(m_i, \kappa_i) - u(\kappa_i)(1 - F_i(\kappa_i))$ and note that $\gamma_i > 0$. As $m_{i^{opt}} \rightarrow \infty$ and $\kappa_{i^{opt}} \rightarrow \infty$, there will eventually come a time where $2\alpha(m_{i^{opt}}, \kappa_{i^{opt}}) + u(\kappa_{i^{opt}})(1 - F_{i^{opt}}(\kappa_{i^{opt}})) < \gamma_i$ for any i . At this point we will have

$$\begin{aligned} UCB_i &\leq LCB_i + 2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) && \text{(Lemma 2)} \\ &= U_{i^{opt}} - \Delta_i + 2\alpha(m_i, \kappa_i) + u(\kappa_i)(1 - F_i(\kappa_i)) && \text{(definition)} \\ &= U_{i^{opt}} - \gamma_i && \text{(definition)} \\ &\leq UCB_{i^{opt}} - \gamma_i && \text{(Lemma 2)} \\ &\leq LCB_{i^{opt}} + 2\alpha(m_{i^{opt}}, \kappa_{i^{opt}}) + u(\kappa_{i^{opt}})(1 - F_{i^{opt}}(\kappa_{i^{opt}})) - \gamma_i && \text{(Lemma 2)} \\ &< LCB_{i^{opt}} && \text{(time reached)} \\ &\leq LCB_{i^*} && \text{(choice of } i^*) \end{aligned}$$

and so i will be eliminated outright.

Now, let i_r^* be the configuration with highest LCB at the end of round r , chosen at Line 19, and let i^{opt} be an optimal configuration. During a clean execution we will have

$$\begin{aligned} UCB_{i^{opt}} &\geq U_{i^{opt}} && \text{(Lemma 2)} \\ &\geq U_{i_r^*} && (i^{opt} \text{ is optimal)} \\ &\geq LCB_{i_r^*} && \text{(Lemma 2)} \end{aligned}$$

so i^{opt} will not be eliminated. On the other hand, by the above, any suboptimal configuration i will eventually be eliminated, so eventually only the optimal configuration(s) will remain.

Finally, suppose that a clean execution of OUP is run until $2\alpha_{i^{opt}} + u(\kappa_{i^{opt}})(1 - F_{i^{opt}}(\kappa_{i^{opt}})) \leq \epsilon$. Then we will have

$$\begin{aligned} U_{i^{opt}} - U_{i^*} &\leq UCB_{i^{opt}} - LCB_{i^*} && \text{(Lemma 2)} \\ &\leq UCB_{i^{opt}} - LCB_{i^{opt}} && \text{(choice of } i^*) \\ &\leq 2\alpha_{i^{opt}} + u(\kappa_{i^{opt}})(1 - F_{i^{opt}}(\kappa_{i^{opt}})) && \text{(Lemma 2)} \\ &\leq \epsilon && \text{(assumption)} \end{aligned}$$

so the returned configuration i^* is ϵ -optimal. \square

Lemma 3: *An execution of COUP is clean for all phases with probability at least $1 - \frac{\delta}{2}$.*

Proof. For any p, i, m_i, κ_i , define the good events:

$$\begin{aligned}\mathcal{G}_{p,i,m_i,\kappa_i}^{(U)} &= \{|U_i(\kappa_i) - \widehat{U}_i(m_i, \kappa_i)| \leq (1 - u(\kappa_i)) \cdot \alpha_p(m_i, \kappa_i)\} \\ \mathcal{G}_{p,i,m_i,\kappa_i}^{(F)} &= \{|F_i(\kappa_i) - \widehat{F}_i(m_i, \kappa_i)| \leq \alpha_p(m_i, \kappa_i)\} \\ \mathcal{G}_{p,i,m_i,\kappa_i} &= \mathcal{G}_{p,i,m_i,\kappa_i}^{(U)} \cap \mathcal{G}_{p,i,m_i,\kappa_i}^{(F)} \\ \mathcal{G}_{p,i} &= \bigcap_{m_i=1}^{\infty} \bigcap_{l_i=1}^{\infty} \mathcal{G}_{p,i,m_i,2^{l_i-1}} \\ \mathcal{G}_p &= \mathcal{G}_{p,1} \cap \dots \cap \mathcal{G}_{p,n_p}.\end{aligned}$$

An execution is clean in phase p if \mathcal{G}_p holds. The probability that an execution is not clean in any phase is

$$\begin{aligned}\Pr\left(\bigcup_{p=1}^{\infty} \overline{\mathcal{G}_p}\right) &\leq \sum_{p=1}^{\infty} \sum_{i=1}^{n_p} \sum_{m_i=1}^{\infty} \sum_{l_i=1}^{\infty} \left(\Pr(\overline{\mathcal{G}_{p,i,m_i,2^{l_i-1}}}^{(F)}) + \Pr(\overline{\mathcal{G}_{p,i,m_i,2^{l_i-1}}}^{(U)})\right) \quad (\text{union bound}) \\ &\leq \sum_{p=1}^{\infty} \sum_{i=1}^{n_p} \sum_{m_i=1}^{\infty} \sum_{l_i=1}^{\infty} \frac{4\delta}{36p^2 n_p m_i^2 l_i^2} \quad (\text{Hoeffding's inequality}) \\ &\leq \frac{\delta}{2} \quad (\text{Basel problem}).\end{aligned}$$

□

Lemma 4: The proof is completely analogous to that of Lemma 2 and is omitted.

Lemma 5: *An execution of COUP is characteristic for all phases with probability at least $1 - \frac{\delta}{2}$.*

Proof. For any p , and for any $i \in \mathcal{A}_p$, from the definition of OPT^{γ_p} , the probability that $U_i \leq OPT^{\gamma_p}$ is at most $1 - \gamma_p$. So the probability that all n_p configurations have $U_i \leq OPT^{\gamma_p}$ is at most $(1 - \gamma_p)^{n_p}$, which is at most $(e^{-\gamma_p})^{n_p} = e^{-n_p \gamma_p} \leq \frac{3\delta}{\pi^2 p^2}$. Summing over $p = 1, 2, \dots$, the probability that in any phase p all n_p configurations have $U_i \leq OPT^{\gamma_p}$ is at most $\frac{\delta}{2}$. □

Theorem 2: *If COUP is run with parameters $(\epsilon_1, \epsilon_2, \dots)$ and $(\gamma_1, \gamma_2, \dots)$ then with probability at least $1 - \delta$ it returns an (ϵ_p, γ_p) -optimal configuration at the end of every phase $p = 1, 2, \dots$*

Proof. By Lemmas 3 and 5 an execution is both clean and characteristic with probability at least $1 - \delta$. Because it is characteristic, there is some i_p with $U_{i_p} \geq OPT^{\gamma_p}$. Let i^* be the configuration returned at the end of phase p during a clean and characteristic execution. We will have

$$\begin{aligned}U_{i^*} &\geq LCB_{i^*} && (\text{Lemma 4}) \\ &\geq \max_{i \in \mathcal{A}_p} UCB_i - \epsilon_p && (\text{phase termination condition}) \\ &\geq UCB_{i_p} - \epsilon_p && (\text{maximum}) \\ &\geq U_{i_p} - \epsilon_p && (\text{Lemma 4}) \\ &\geq OPT^{\gamma_p} - \epsilon_p && (\text{characteristic}).\end{aligned}$$

□

B IMPROVED CAPTIME DOUBLING CONDITION

We discuss the change in doubling condition as applied to OUP, but the same argument holds for COUP more generally on a per-phase basis, as well as for UP. Line 9 of Algorithm 1 shows the doubling condition introduced by UP:

9: **if** $2\alpha(m_i, \kappa_i) \leq u(\kappa_i)(1 - \widehat{F}_i)$ **then** \triangleright captime doubling condition

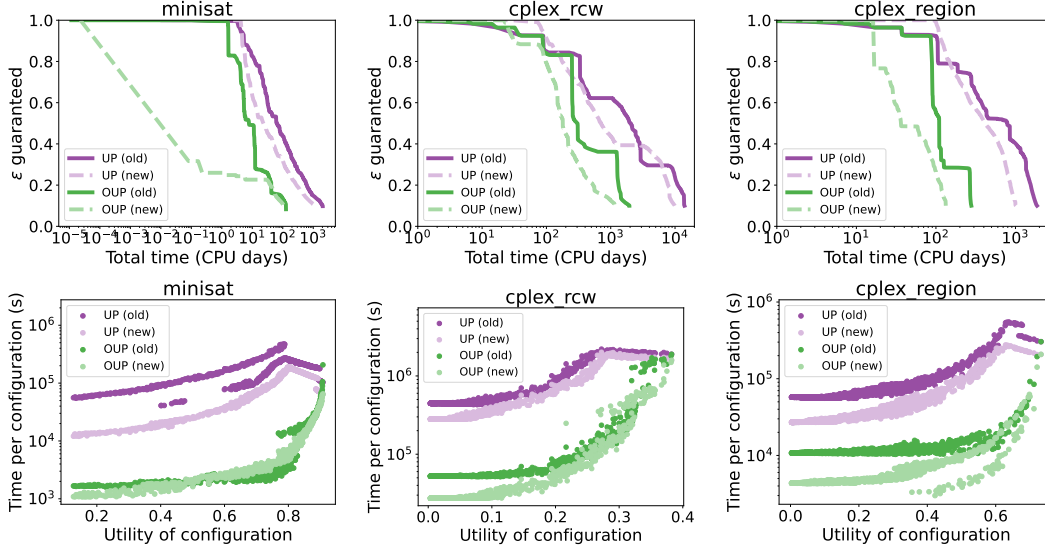


Figure 6: Improvement of the new doubling condition on UP and OUP, using the log-Laplace utility function. Other utility functions show similar improvement.

The term $2\alpha(m_i, \kappa_i)$ can be interpreted as the error incurred due to sampling because this is the term that comes from taking samples and applying the statistical bound of Hoeffding’s inequality. The term $u(\kappa_i)(1 - \hat{F}_i)$ can be interpreted as the error incurred due to capping because runs longer than κ_i can have a utility of at most $u(\kappa_i)$, and there will only be (approximately) $1 - \hat{F}_i$ of these. Balancing these two terms seems like a sensible approach, but a more careful analysis suggest something slightly different.

To emphasize that our uncertainty is the result of both sampling and capping, we re-write the upper and lower confidence bounds without cancelling terms as

$$\begin{aligned} UCB_i &= \hat{U}_i + (1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i) + 0, \\ LCB_i &= \hat{U}_i - (1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i) - u(\kappa_i)(1 - \hat{F}_i + \alpha(m_i, \kappa_i)). \end{aligned}$$

It can now be seen more clearly that the confidence width decomposes into a sum of two terms. The first represents uncertainty about runs from below the captme κ_i , and the second represents uncertainty about runs from above κ_i :

$$UCB_i - LCB_i = \underbrace{2(1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i)}_{\text{error from runs below } \kappa_i} + \underbrace{u(\kappa_i)(1 - \hat{F}_i + \alpha(m_i, \kappa_i))}_{\text{error from runs above } \kappa_i}.$$

Because i was chosen at line 7 to have maximum UCB , we have $UCB_i - LCB_i \geq UCB_{i_{opt}} - LCB_i \geq \Delta_i$, so shrinking the confidence width $UCB_i - LCB_i$ is necessary to prove a better optimality guarantee for configuration i . But because the confidence width is a sum of two terms, shrinking it toward zero requires shrinking both terms toward zero; the confidence width is dominated by the largest term. The new captme doubling condition we propose shrinks these two terms in a balanced way:

\mathcal{G}' : **if** $2(1 - u(\kappa_i)) \cdot \alpha(m_i, \kappa_i) \leq u(\kappa_i)(1 - \hat{F}_i + \alpha(m_i, \kappa_i))$ **then** \triangleright captme doubling condition

Empirically, this is generally an improvement over the old doubling condition and sometimes a large improvement. We set $\epsilon = \delta = .1$ and run both OUP and UP using both the old and new doubling conditions. Figure 6 clearly shows the improvement to both procedures in terms of the optimality guarantee being made and the time spent per configuration.

C DESCRIPTION OF PERCENTAGE-GAP EXPERIMENT

In Section 5.3 we compare COUP and OUP to the following procedures.




Figure 7: Average configurator performance as measured by the percentage gap to best configuration in the dataset optimizing an expected runtime. Quantile parameter $\delta = .2$

Successive Halving (Jamieson & Talwalkar, 2016) is a simple procedure that “successively halves” the number of configurations based on their performance until only one remains. It has a budget which it allocates equally to surviving configurations. We varied the budget between 22000 and 186000 for the `minisat` dataset and between 22000 and 350000 for the CPLEX datasets, which is the largest range possible based on the numbers of configurations and instances.

Hyperband (Li et al., 2018) builds on Successive Halving, allocating the budget more effectively across configurations. We set Hyperband to optimize the utility function while minimizing runtime as its resource budget. We set the multiplier $\eta \in \{2, 3, 4, 5, 6\}$ so that there would be about five brackets, as recommended by the authors, and then set the resource parameter R to be as high as the dataset would permit. We ran Hyperband and Successive Halving with runs capped at $\kappa \in \{1, 10, 100, 500\}$ for the `minisat` dataset and $\kappa \in \{10, 100, 1000, 5000\}$ for the CPLEX datasets.

AC-Band Brandt et al. (2023) builds further on Hyperband, but is specifically designed for algorithm configuration. We set AC-Band to optimize the utility function, but we found this did not work very well, so we also reported its performance when optimizing its default metric, which is better. We suspect that this utility function is uniquely suited to AC-Band. AC-Band runs a number of configurations in parallel and terminates them all as soon as one finishes. Meanwhile, its default evaluation metric counts the number of times a configuration finishes first. The AC-Band paper reports the percentage gap from best average runtime, though it optimizes a different function. For consistency, we report this percentage gap. Also, we fixed what we believe to be a bug in the code from the AC-Band repo. Line 44 of the `cse.py` file calls the `env.run` function with a hard-coded captime value of 900. This is the maximum captime of the smaller, `minisat` dataset, but the larger CPLEX datasets were collected with a maximum captime of 10000 seconds. We used a captime of 10000 seconds for the CPLEX datasets (though this made only a little difference).

ImpatientCapsAndRuns (ICAR) (Weisz et al., 2020) builds on a different line of research, and is designed to discard poorly-performing configurations quickly. It measures performance according to an expected capped runtime, where the cap is set so that only a δ -fraction of runs time out. Following the experiments reported in their paper, we set $\epsilon = 0.25$ for the `minisat` dataset and $\epsilon = 0.1$ for the CPLEX ones, which were the smallest values we could use without running out of instances. We used $\delta \in \{0.1, 0.2\}$ and varied $\gamma \in \{0.01, 0.02, 0.05\}$ (note that these parameters have different meaning than in our paper). We varied the boolean parameters with the same configurations as the authors.

Structured Procrastination with Confidence (SPC) (Kleinberg et al., 2019) makes the same optimality guarantee as ICAR. We ran SPC with a θ -multiplier parameter of 2 and 3, and checked results after 0.1, 1, 10 and 100 CPU days of total compute time for the `minisat` dataset and 1, 10, 100 and 1000 CPU days for the CPLEX datasets. The total compute time implies a certain ϵ , depending on the dataset. Since they optimize the same objective function, we used the same δ values to measure the performance of SPC as we used for ICAR.