

Practical, Utilitarian Algorithm Configuration

Devon R. Graham¹, Eros Rojas Velez¹, Kevin Leyton-Brown¹

¹Department of Computer Science
University of British Columbia

drgraham@cs.ubc.ca, erojas20@student.ubc.ca, kevinlb@cs.ubc.ca

Abstract

Utilitarian algorithm configuration identifies a parameter setting for a given algorithm that maximizes a user’s utility. Utility functions offer a theoretically well-grounded approach to optimizing decision-making under uncertainty and are flexible enough to capture a user’s preferences over algorithm runtimes (e.g., they can describe a sharp cutoff after which a solution is no longer required, a per-hour cost for compute, or diminishing returns from algorithms that take longer to run). COUP is a recently-introduced utilitarian algorithm configuration procedure which was designed mainly to offer strong theoretical guarantees about the quality of the configuration it returns, with less attention paid to its practical performance. This paper closes that gap, bringing theoretically-grounded, utilitarian algorithm configuration to the point where it is competitive with widely used, heuristic configuration procedures that offer no performance guarantees. We present a series of improvements to COUP that improve its empirical performance without degrading its theoretical guarantees and demonstrate their benefit experimentally. Using a case study, we also illustrate ways of exploring the robustness of a given solution to the algorithm selection problem to variations in the utility function.

Code — <https://github.com/drgrhm/practical-coup>

Extended version — <https://arxiv.org/abs/2510.14683>

1 Introduction

Algorithm configuration is the process of finding a complete setting of an algorithm’s parameters that lead it to perform well on a given set of input instances. Often, the algorithm of interest is a solver for a hard computational problem such as Boolean Satisfiability (SAT). Since SAT is NP-complete, we are unlikely to be able to find a single solver or configuration that performs well on all possible inputs. In practice, however, we are never interested in solving *all possible* inputs. Instead, we tend to be interested in solving input instances drawn from a given distribution.

Algorithm configuration is thus a bandit problem (Lattimore and Szepesvári 2020) where each configuration is an arm, and running a configuration on an input instance means drawing a sample from that arm’s underlying distribution. Unlike the standard bandit setting, in which there is a unit cost

associated with pulling an arm, in algorithm configuration the cost is the time it takes for the algorithm to solve the input. We can also choose whether each such evaluation should be censored (“capped”) at some maximum runtime. The objective in algorithm configuration is not regret minimization, but rather (approximate) best arm identification.

Algorithm configuration procedures have traditionally aimed to minimize the expected runtime of the returned configuration. The most practical such methods make these decisions according to carefully chosen heuristics (Birattari et al. 2002; Hutter et al. 2009; Ansótegui, Sellmann, and Tierney 2009; Hutter, Hoos, and Leyton-Brown 2011; López-Ibáñez et al. 2016; Lindauer et al. 2022). A second line of work has identified (considerably less practical) algorithms that offer worst-case guarantees about the runtime required to identify approximately optimal solutions (Kleinberg, Leyton-Brown, and Lucier 2017; Weisz, Gyorgy, and Szepesvári 2018; Kleinberg et al. 2019; Weisz, Gyorgy, and Szepesvári 2019; Weisz et al. 2020). For the purposes of this paper, we define an optimality guarantee as depending on three parameters: ϵ , γ , and δ . With failure probability δ , we want to prove both that the returned configuration is within an additive factor of ϵ from the best configuration sampled, and that the probability of sampling a configuration that is better than the best one sampled so far is at most γ .

It has recently been argued that algorithm configuration should rest on the firmer decision-theoretic foundation of maximizing a utility function over algorithm runtime (Graham, Leyton-Brown, and Roughgarden 2023a). That is, let u be a function that maps runtimes to utility values in $[0, 1]$, encoding the user’s preferences about runtimes of different lengths; we should then maximize u ’s expected value. A recent line of work has proposed procedures that offer worst-case guarantees about the runtime required to approximately optimize a utilitarian objective (Graham, Leyton-Brown, and Roughgarden 2023b; Brandt et al. 2023). The most recent such procedure is COUP (Continuous, Optimistic Utilitarian Procrastination; Graham and Leyton-Brown 2024), which was shown empirically to perform favorably with respect to other configuration procedures that offer (in some cases, somewhat incomparable) optimality guarantees: Successive Halving (Jamieson and Talwalkar 2016), Hyperband (Li et al. 2018), AC-Band (Brandt et al. 2023), ImpatientCapsAndRuns (Weisz et al. 2020) and Structured Procrastination with

Confidence (Kleinberg et al. 2019). Pseudocode for COUP is included in Appendix A of the extended version of the paper, but in essence it is an instantiation of the UCB algorithm (Lai and Robbins 1985; Lai 1987), coupled with a capttime exploration technique and a condition for expanding the set of configurations under consideration. For each configuration i , COUP maintains upper and lower confidence bounds UCB_i and LCB_i . In each iteration it chooses the configuration $i' = \arg \max_i UCB_i$ and performs a run of it. In this way it works to improve $\epsilon^* = \max_i UCB_i - \max_i LCB_i$, which is the best optimality guarantee it can make. If COUP detects that the capttime it is using is too small, it will double it. If COUP detects that its set of configurations is too small, it will add more. COUP is *anytime* in its guarantees with respect to both ϵ and γ : as it runs it makes tighter and tighter optimality guarantees with respect to larger and larger sets of configurations. COUP is also *adaptive* to the inputs it is given: it provably makes better guarantees in less time when given easier inputs.

This paper begins in Section 2 by describing a series of improvements to COUP that boost its empirical performance without compromising its theoretical guarantees. First, we refine the confidence bounds used by COUP, making the optimization process more efficient. Second, we incorporate a different version of the UCB algorithm that specifically targets the best arm identification problem. Third, we suggest a new condition for when COUP should add new configurations, which frees the user from having to provide a set of parameters ahead of time, instead adapting to the inputs it sees. And fourth, we make COUP model-based, using a boosted forest of regression trees to help COUP find new configurations. In Section 3 we show for the first time that an algorithm configuration method offering theoretical guarantees (our modified version of COUP) robustly achieves performance roughly equivalent to that of SMAC (Hutter, Hoos, and Leyton-Brown 2011; Lindauer et al. 2022), a heuristic procedure widely seen as the empirical state of the art. Finally, in Section 4, we address a shortcoming of the utilitarian framework: that it is not always easy for the end user to choose the exact utility function to optimize. Via a case study on data from one of the most prominent international evaluations of heuristic algorithms, the SAT Competition, we illustrate how runtime data can be used to reason about the utility functions for which a given algorithm would be considered optimal.

2 Practical Improvements to COUP

We now present four of practical improvements to COUP that significantly improve its performance, both in terms of the optimality guarantees it is able to make and in terms of the empirical quality of the configuration it returns. Section 2.1 discusses a tightening of the confidence bounds used by COUP, Section 2.2 argues for using a variant of the UCB algorithm, Section 2.3 discusses an improvement to the way COUP incorporates new configurations during its search process, and Section 2.4 discusses using a model to help guide COUP’s search for new configurations.

2.1 Improved Confidence Bounds

A key technical tool used in COUP is Hoeffding’s Inequality (Hoeffding 1963). This concentration inequality asserts that for bounded random variables, the probability that the observed sample mean is far from the true population mean is exponentially small in the number of samples taken. However, Hoeffding’s Inequality involves a key approximation, which is tight when the true mean is close to 1/2 (in the case of random variables bounded on $[0,1]$), but grows worse when the true mean is close to either 0 or 1—a situation encountered often in utilitarian algorithm configuration. We can do away with this approximation by appealing directly to the “Chernoff–Hoeffding Lemma”, and solving numerically for the confidence bounds, which are then guaranteed to be tighter than a naive application of Hoeffding’s Inequality; for a detailed exposition see Lattimore and Szepesvári (2020), Chapter 10. The confidence bounds come from solving an optimization problem that bounds the KL-Divergence between the observed sample mean and the bounding quantity. We thus refer to them as “KL bounds”.

COUP maintains a capttime κ , which is the maximum time it will spend on each run of a configuration. To derive upper and lower confidence bounds for a configuration i ’s mean utility, COUP needs to estimate its capped mean runtime $U_i = \mathbb{E}_{t_i}[u(\min(t_i, \kappa))]$ and its probability of not capping $F_i = \Pr_{t_i}(t_i \leq \kappa)$. Both of these estimates can be improved by switching from Hoeffding bounds to KL bounds. Define $d(p, q) = p \log(p/q) + (1-p) \log((1-p)/(1-q))$. Suppose n is the current number of configurations, m is the number of samples we have taken for configuration i , and κ is the capttime we have taken them at. If \hat{U}_i is the empirical capped average runtime, \hat{F}_i is the empirical fraction of completed runs, and $a = \frac{1}{m} \ln \frac{27n^2 m^2 (\log(\kappa+1))^2}{\delta}$ then

$$UCB_i = \max \{u \in [0, 1] : d(\hat{U}_i, u) \leq a\}, \text{ and}$$

$$LCB_i = \min \{u \in [0, 1] : d(\hat{U}_i, u) \leq a\}$$

$$- u(\kappa)(1 - \min \{u \in [0, 1] : d(\hat{F}_i, u) \leq a\})$$

are confidence bounds that hold for all configurations at all times during COUP’s execution with probability at least $1 - \delta$ (detailed derivations can be found in Appendix B of the extended version of the paper). This alteration essentially makes COUP an implementation of the KL-UCB algorithm (Cappé et al. 2013) for bandit problems. In Figure 2 (Left) in Section 3 we show that COUP performs better when using these refined bounds.

2.2 Lower Upper Confidence Bound

The UCB algorithm is a natural choice as a framework for algorithm configuration procedures. Indeed, Graham and Leyton-Brown (2024) argue that UCB is the best way to apply utilitarian algorithm configuration to infinite parameter spaces, and show that using UCB is a significant improvement over configuration procedures that implement other bandit procedures. However, UCB is essentially designed for regret minimization, while the algorithm configuration problem is a case of best arm identification. Thankfully, the

LUCB (*Lower Upper Confidence Bound*) algorithm (Gabilon, Ghavamzadeh, and Lazaric 2012; Kalyanakrishnan et al. 2012) is a variant of UCB that was designed to be applied to the best arm identification problem. Actually, the LUCB algorithm was designed to solve the m -best arm identification problem, where a learner is tasked with identifying the top m arms from a set of n arms. Although we are only concerned with the special case of $m = 1$, it is easier to understand what LUCB is doing by considering the more general case. In each iteration, LUCB selects and samples *two* arms (i.e., configurations). First it samples the arm with the smallest lower confidence bound from the m empirically best arms. Next it samples the arm with the largest upper confidence bound from among the $n - m$ remaining arms. This way, the LUCB algorithm attempts to tease apart the m best arms from the $n - m$ worst arms. Figure 2 (Left) in Section 3 shows an example of the improvement achieved by using the LUCB algorithm. This is drastic early in a search.

2.3 Adding New Configurations

COUP runs in phases. In each phase it reduces both ϵ and γ by an amount that is specified by the user at runtime. This approach is non-adaptive; it requires that the user specify a sequence of ϵ and γ values as parameters, which is not only burdensome but also does not automatically take into account any properties of the actual runtime data COUP sees. For instance, if COUP finds that the best utility it can prove is relatively low, then it should tend to favor sampling new configurations more than improving the optimality guarantee of existing configurations. To this end, we define two conditions for adding new configurations that adapt to the relative sizes of the configurations' confidence bounds.

The first condition simply balances the ϵ and γ that COUP can prove. Defining $i^* = \arg \max_i LCB_i$ and $i' = \arg \max UCB_i$, if COUP adds new configurations when $UCB_{i'} - LCB_{i^*} < \gamma$ then it will work to reduce the ϵ it can guarantee, until this is smaller than the γ it can guarantee, at which point it adds a new configuration. This frees the user from having to specify a sequence of ϵ values, but does not have the property that it explores more new configurations when the best utility it can prove is relatively low.

The second condition does have this property. At any point during its execution, COUP will recommend the configuration $i^* = \arg \max LCB_i$, the configuration with largest lower confidence bound. This is the configuration about which it can prove the best guarantee. The quantity LCB_{i^*} is the highest expected utility that COUP can guarantee at any time. We would like to push LCB_{i^*} as close to 1 as possible because then we will have (i) a very good configuration, and (ii) a very strong guarantee about this configuration. So we can think of the quantity $1 - LCB_{i^*}$ as the utility "still on the table"; the largest amount of utility by which we might hope to improve our guarantee. We can write $1 - LCB_{i^*} = 1 - UCB_{i'} + UCB_{i'} - LCB_{i^*}$. The term $UCB_{i'} - LCB_{i^*}$ is the best ϵ we can guarantee with respect to our existing set of configurations; the maximum amount by which we can shrink $1 - LCB_{i^*}$ by running existing configurations. On the other hand, since $UCB_{i'}$ is the largest upper bound from our current sample of configurations, we

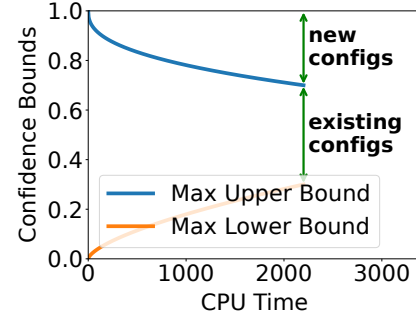


Figure 1: Illustration of the decomposition that leads to our adaptive strategy for incorporating new configurations.

can be confident that the configurations we have seen so far have utility at most $UCB_{i'}$. So the term $1 - UCB_{i'}$ is the utility potentially to be gained from other configurations that we have not yet seen. This is illustrated in Figure 1.

The magnitude of $UCB_{i'} - LCB_{i^*}$ tells us what is potentially to be gained by running our existing configurations, while the magnitude of $1 - UCB_{i'}$ tells us what is potentially to be gained by sampling and running new configurations. This suggests a simple condition: to add new configurations when $UCB_{i'} - LCB_{i^*} < 1 - UCB_{i'}$, which means to add new configurations when there is more to be gained from doing so than not. However, with this condition there is a slight problem. If we add new configurations when $UCB_{i'} - LCB_{i^*} < 1 - UCB_{i'}$, then while COUP is running existing configurations it has $UCB_{i'} - LCB_{i^*} \geq 1 - UCB_{i'}$. This means that we will always have $UCB_{i'} \geq \frac{1+LCB_{i^*}}{2}$ and so $UCB_{i'} - LCB_{i^*} \geq \frac{1-LCB_{i^*}}{2} \geq \frac{1-U_{i^*}}{2}$. This means that the ϵ we can guarantee will never approach 0 but will asymptote (unless $U_{i^*} = 1$, in which case we will have $UCB_{i^*} = 1$ also and we will not sample any more configurations, which is what we want since U_{i^*} is already maximal).

Both conditions have downsides but by combining them together we get a condition that avoids these drawbacks. If we add new configurations when $(UCB_{i'} - LCB_{i^*})^2 < \gamma(1 - UCB_{i'})$, then we will tend to add new configurations when $1 - UCB_{i'}$ is large, but then this will cause γ to be reduced and COUP will tend to make a corresponding reduction to $\epsilon^* = UCB_{i'} - LCB_{i^*}$. Thus this condition will tend to continue to shrink ϵ and also is adaptive to the inputs given. We note that this condition is equivalent to comparing ϵ^* on the left to the geometric mean of γ and $1 - UCB_{i'}$ on the right.

2.4 Model-guided Configuration Search

COUP can be improved by incorporating a learning model that predicts the performance of candidate configurations before spending significant amounts of time investigating them. If good configurations are rare enough, then we will be all but guaranteed to never encounter them by random sampling alone. A model will at least have the possibility of finding such configurations, and there is evidence for this: Hutter, Hoos, and Leyton-Brown (2011) report that SMAC outperforms ROAR, its model-free variant, in all scenarios and

substantially outperforms ROAR on some large configuration spaces, indicating that at least in some cases the prediction model is able to find significantly better configurations than random sampling alone.

As COUP runs, it generates the performance data needed to train a performance model and we can use the trained model to help identify promising areas of the search space to investigate. Areas that are predicted to yield configurations with high performance will be investigated more, as will areas that have high uncertainty. We use a gradient-boosted set of regression trees as our base learning model to predict configuration performance. We use the XGBoost (Chen and Guestrin 2016) implementation for gradient boosting. In order to produce uncertainty estimates we perform a bootstrap aggregation procedure, sampling subsets of configurations and training an XGBoost model on each one. For each candidate configuration we make a prediction with each of the models, giving us a distribution of values. We then consider 95% confidence intervals, giving us an uncertainty band around each prediction. See Appendix E in the extended version of the paper for more details.

Given a trained model, we employ the same procedure used by SMAC (Hutter, Hoos, and Leyton-Brown 2011) to identify promising configurations to investigate. Briefly, we start with the ten best configurations seen so far and perform a local search starting with each. For discrete parameters, we consider the set of all configurations that differ in exactly one parameter value. For continuous parameters, we normalize values to the range $[0, 1]$ and sample four neighbors according to a truncated Gaussian distribution with mean equal to the parameter value and standard deviation equal to 0.2. We continue this search as long as the neighbors have a larger upper confidence bound. Once we have sampled 10 neighbors, we use the model to predict their performance, as well as the performance of 10,000 randomly selected configurations, choosing the one that has the highest predicted 95% upper confidence bound as the next candidate.

The parameter γ described above indicates how likely it is that we sample a configuration that is better than any we have seen so far. As we drive γ to 0, we grow more and more sure that we are unlikely to find a better configuration than our current incumbent. However, making this proof requires a growing number of randomly sampled configurations. To continue making this guarantee, we only sample every second configuration using the model, with the other half of the samples chosen randomly. In this case we have a sample of $n' = n_0 + \lfloor \frac{n-n_0}{2} \rfloor \geq \frac{n}{2}$ random configurations, where n_0 is the number of initial configurations considered, before the model is trained. Thus, we need to sample twice as many configurations to make the same γ guarantee, but this is just a (small) constant, and γ is still driven toward 0 as $n \rightarrow \infty$ (see Appendix C of the paper’s extended version for a full derivation). We retrain the model before asking it to suggest configurations, thus incorporating the new training data we have observed. In Section 3 we show empirically that using a prediction model helped COUP find good configurations more quickly, and that in general configurations found by the model tended to be better than the configurations found by random sampling.

3 COUP’s Empirical Performance

We have argued theoretically, and provide experimental evidence below, that our changes from Section 2 only improved COUP, first by tightening the bounds and refining the implementation of UCB (Figure 2, Left) and second by incorporating a performance model (Figure 2, Center). We also compare COUP to the heuristic procedure that is most widely used in practice: SMAC. While SMAC was initially proposed for the optimization of expected runtime, it is able to optimize any objective function, and so can be straightforwardly extended to a utilitarian objective. We thus compared COUP’s performance to that of SMAC 3 (Lindauer et al. 2022) in scenarios consisting of four different SAT solvers and 20 different instance sets.

We perform experiments using the ACLib library of algorithm configuration benchmarks (Hutter et al. 2014). ACLib is well-established as a benchmark set for demonstrating algorithm configuration performance. To the best of our knowledge, ours is the first paper to demonstrate utilitarian algorithm configuration with general optimality guarantees in an online setting, using real runs from real solvers. In contrast, previous work has used precomputed or synthetic data.

We ran experiments on a 14-node CPU cluster. Each node consisted of 32 dual-socket Intel Xeon E5-2683 CPUs with 96GB of memory. Each COUP scenario (solver and dataset combination) was trained for 96 hours with 12GB of memory per job, with each validation instance having a timeout of 60 seconds once training had been completed. We use the *runsolver* timing tool (Roussel 2011) to obtain runtime measurements that accurately reflected the total CPU time consumed by each solver.

For each solver–instance set pair, we evaluated performance on a random sample of 2000 training instances (unless the dataset consisted of fewer than 2000 instances). After each iteration of COUP’s training, we saved the best configuration and validated each unique configuration found. We then evaluated SMAC’s incumbent configuration at the same points in time and on the same set of validation instances. We limited the individual solver runs that SMAC performed to 60 seconds and used the same limit when validating both COUP and SMAC. We used an initial captime of 1 second for COUP, which is the smallest execution time that *runsolver* is able to enforce. We set COUP’s failure probability $\delta = 0.01$.

The utility function we used has two parameters: κ_0 and κ_1 . If $t < \kappa_0$ then $u(t) = 1$, indicating that all runs faster than κ_0 are identical to us and essentially instantaneous. If $t > \kappa_1$ then $u(t) = 0$, indicating that all runs longer than κ_1 are essentially worthless to us. Between these two extremes, when $t \in [\kappa_0, \kappa_1]$ we have $u(t) = \frac{\log(t/\kappa_1)}{\log(\kappa_0/\kappa_1)}$, meaning that if runtime is doubled then utility is reduced by a constant factor for any runtime in $[\kappa_0, \kappa_1]$. Optimizing this utility function, which amounts to a scaled and shifted logarithmic transform of runtime, will return configurations that are good for runtimes that are of different orders of magnitude, making it applicable over a broad range of problems. Moreover, using a logarithmic transform of runtime has been shown to improve the quality of algorithm runtime prediction models (Xu et al. 2008). We set $\kappa_0 = 0.001$ seconds, which is approximately

the time it takes the operating system to launch a process and thus is the absolute minimum runtime we could observe. We set $\kappa_1 = 1$ hour, which is the maximum amount of time we imagine waiting for any one run to finish. Of course other utility functions are possible. We note that some utility functions might be more “strict” than others and will take more time to optimize. In general, the utility functions we consider drop from 1 to 0 as runtime increases. If this drop tends to happen at large runtime values, then COUP will tend to have to do longer runs, and take more time (depending on the given runtime distributions).

In Section 2.1 we argued that using the KL bounds instead of naive Hoeffding bounds means that COUP is able to prove a better ϵ more quickly. In section Section 2.2 we argued that using the LUCB algorithm rather than the vanilla UCB algorithm also means COUP proves better guarantees more quickly. Figure 2 (Left) shows empirically that both of these changes improve the ϵ that COUP is able to prove. This plot shows results for 100 configurations of the Spear SAT solver on the FACTORING instance set from ACLib. Other scenarios show similar results. We have tended to observe that using the improved bound helped the most in terms of returning a configuration with good utility, while the use of the LUCB algorithm helped the most in terms of guaranteeing a good ϵ .

In Section 2.4 we argued that using a prediction model would help COUP find better configurations overall. Figure 2 (Center) evaluates this process. We first sampled 30 configurations at random, then alternated between sampling at random and sampling according to the model. Configurations with a higher index (x -axis) were sampled later than configurations with a lower index. Not only did configurations sampled according to the model tend to be better than configurations sampled at random, but we can clearly see that the quality of the configurations suggested by the model improved with time, as the model received more training data. The plot shows the performance of the Spear solver on the FACTORING instance set from ACLib. Other scenarios showed qualitatively similar results.

Our augmented version of COUP identified good configurations about as quickly as SMAC at our cutoff time of four wall-clock days. For every (solver, instance) pair, we define the normalized utility as the utility achieved by a given algorithm configuration run divided by the best such utility observed across both the SMAC and COUP runs. Figure 2 (Right) is an empirical CDF of the difference between the normalized utilities achieved by COUP and SMAC across the different solver and instance set pairs. We call runs that achieved solutions within 10% of each other effectively equivalent, represented as the region between the dashed red lines. 71 out of our 80 runs fell within this region; of what remains, 5 had a difference less than -0.1 (meaning SMAC outperformed), and 4 had a difference greater than 0.1 (COUP outperformed). We performed the same analysis at various earlier cutoff times and observed qualitatively similar results.

It is worth noting that COUP solves a considerably harder problem than SMAC. Rather than simply identifying strong algorithm configurations, it also proves anytime theoretical guarantees about the quality of the solution it has found at

every point in time, both in terms of the amount by which the returned configuration might be worse than some other configuration under consideration (ϵ) and in terms of the probability that a better configuration could be sampled (γ). Figure 3 shows the optimality guarantees proved by COUP for different solver–instance set pairs. Generally, both ϵ (Left) and γ (Right) are driven toward 0, but some trends are evident. ProbSAT struggles to prove a good ϵ (purple curves tend to be higher). All solvers struggle to prove a good ϵ on at least one dataset (dark cluster at the top of the plot), and could use with more training time to push ϵ even farther down. Figure 3 (Right) begins with a solid line (hence the compressed x axis) because COUP tightens only ϵ until it has learned sufficiently about its initial set of 100 configurations. Plots for individual runs can be found in Appendix G of the extended version of the paper.

Finally, Graham and Leyton-Brown (2024) included a comparison of basic COUP to SuccessiveHalving (Jamieson and Talwalkar 2016), Hyperband (Li et al. 2018), AC-Band (Brandt et al. 2023), Structured Procrastination (Kleinberg et al. 2019), and ImpatientCapsAndRuns (Weisz et al. 2020). All of these procedures make theoretical guarantees of some sort, but it is important to note that these guarantees are not necessarily comparable to one another. COUP was found to perform well compared to all of these competitors, as was its finite-configuration-set variant, OUP. We included our improved variant of COUP (and OUP) in this comparison. ?? in Appendix D of the paper’s extended version plots the gap between the best configuration found and the best configuration overall for each of these procedures. It can be seen that our improved variants perform even better than their basic counterparts, in some cases quite significantly.

4 Utility Functions from Runtime CDFs

Like all previous work on utilitarian algorithm configuration, we have so far assumed that we know the utility function u being optimized. In practice, however, we may not be sure how to instantiate certain parameters of u . In this section we illustrate a way around such uncertainty: gathering runtime data from which we can draw conclusions and then testing how sensitive our conclusions are to the parameter values.

We illustrate several different approaches to such an analysis by working through a concrete example: the International SAT Competition (<https://satcompetition.github.io>). This is a prominent annual event held since at least 2002. Participants in the competition are asked to submit solvers which will be evaluated on an unknown set of SAT problem instances. Typically, 300 to 600 benchmark instances are used, and each solver is given a maximum of 5000 seconds to solve each instance. Once the submitted solvers have been run on the benchmark instances, we can derive CDFs for each one. We can infer a great deal from these CDFs about which solver is best under various criteria.

The simplest inference we can make is based on First Order Stochastic Dominance (FOSD). For two solvers A and B with runtime CDFs F_A and F_B , we say that A is FOSD over B if $F_A(t) \geq F_B(t)$ for all t in the support of A and B , and $F_A(t) > F_B(t)$ for some t . (The inequality is reversed from the definition usually seen because we are




Figure 2: Left: Suboptimality guarantee provided by COUP, demonstrating the improvements achieved by using the KL bounds and the LUCB algorithm. Center: Lower confidence bounds for configurations sampled by COUP. Configurations sampled according to the model tend to be provably better and to improve with time as more training data is gathered. Right: The CDF of the difference in normalized final utility between COUP and SMAC. Most of the time, COUP and SMAC perform equivalently.




Figure 3: Unlike SMAC, COUP also offers theoretical guarantees. In general, both ϵ and γ are driven toward 0.




Figure 4: The winning SAT solver is FOSD over the 16th ranked solver.




Figure 5: The winning SAT solver is FOSD over the 21st ranked solver for runtimes $> \sim 10$ s.

dealing with runtime, which we want to minimize, instead of money or some other good that we want to maximize.) An immediate result of this definition is that a solver A is FOSD over a solver B if and only if A is expected-utility preferred over B for *any* utility function. For example, consider the runtime data from the 2023 SAT Competition. Only two solvers, the 8th ranked “PReLearn-kissat-PReLearn-kissat” and the 16th ranked “PReLearn-kissat-PReLearn-tern-kissat” are first-order stochastically dominated by the winner (Figure 4). These solvers will never outrank the winner, no matter what evaluation criterion we use. Other solvers are “approximately” first-order stochastically dominated by the winner. For instance, the 21st ranked solver “Kissat.MAB.Rephases” is shown in Figure 5, and the winning solver is FOSD over it for runtimes

greater than about 10 seconds. (For context, fewer than 10% of all solves by all solvers took less than 10 seconds.)

If the organizers decide that all runs less than 10 seconds are essentially equivalent, then this will be reflected in the chosen utility function, which will be constant up to (at least) 10 seconds. In this case we are essentially back to having FOSD. The winning solver is expected-utility preferred over solver 21 for any utility function that is constant up to 10 seconds. For instance one of the simplest evaluation criteria we could use in this scenario is the proportion of instances solved. This amounts to optimizing a step function utility of the form $u(t) = 1$ if $t < \kappa$ and $u(t) = 0$ otherwise, where κ is a parameter. For any $\kappa > 10$ seconds, the winning solver is FOSD over solver 21. In general, a utility function




Figure 6: Left: Regret using the $PAR(2, \kappa)$ utility function. Center: Regret using the $PAR(c, 5000)$ utility function. Right: Regret using the exponential utility function.

decreases from 1 to 0 as t increases. The regions where the utility function decreases most will be the most important regions. If all of the decrease happens in a region where solver A is FOSS over solver B , then A will be expected utility preferred over B , regardless of the precise form of the decrease. We expand on this idea in Appendix F of the paper’s extended version.

The organizers have often adopted a “penalized average runtime” (PAR) scoring function. The score for solving an instance is the time the solver took to solve it, while the score for a timeout is a constant multiplied by the timeout value. This scoring function has two parameters: a penalty factor c , and a maximum timeout κ . We can map this to a utility function in the range $[0, 1]$ by setting $u(t; c, \kappa) = 1 - \frac{t}{c\kappa}$ if $t < \kappa$ and $u(t; c, \kappa) = 0$ otherwise. Ranking solvers according to the $PAR(c, \kappa)$ score is equivalent to (i.e., will always give equivalent rankings as) ranking solvers according to the utility function $u(t; c, \kappa)$.

In recent years, the SAT Competition has evaluated solvers according to a $PAR(c, \kappa)$ utility function with $c = 2$ and $\kappa = 5000$. We might naturally ask how important the choice of values of these parameters is to the final rankings of solvers, since the values $c = 2$ and $\kappa = 5000$ are essentially arbitrary. Is the winning solver only the winner because of these specific values? Might the final rankings have been significantly different for different values? If the rankings change significantly with small changes in the parameter values, then we might be less convinced that the winning solver is really the “best” one, while if the winning solver remains highly ranked for a wide range of parameter values, we would be more confident about this conclusion.

We can plot the utility of each solver as a function of the parameter(s) of the utility function. We can then look at which solver is best for each value of the parameter(s). Figure 6 (Left) and Figure 6 (Center) show the regret for each solver, which is the difference between its utility and the utility of the best solver for that parameter value. Holding $c = 2$ fixed, for κ above about 1000 the winning solver is best, but for smaller κ it is actually the 4th ranked solver (“Kissat_MAB_prop-no.sym”) that is best. For a small range of κ around about 1000 the 3rd ranked solver (“Kissat_MAB_prop”) is best, but all three solvers have very similar performance at this point. If we are trying to choose a solver for some specific application, but we do not know exactly how much time we will be given to solve an instance, we can choose our solver

based on whether we think we will be given more or less than 1000 seconds. However, if we are unsure of whether we have more or less than 1000 seconds, we could choose the 4th ranked solver and our regret would never be too great, whereas if we chose the winning solver our regret could be relatively large. Finally, we can see that when $\kappa = 5000$, the best ranked solver is best for all values of c . This sort of reasoning can be helpful when we have uncertainty about our utility function’s parameters.

We can also consider utility function other than $PAR(c, \kappa)$. For example, we could use an exponential function: $u(t) = e^{-\lambda t}$, where λ is a parameter. If we know the exact value of λ that we want to use, then we can just compute the mean utility for each solver and find which one is best. Otherwise we can plot the mean utility of each solver as a function of λ and see how the ranking changes, as in Figure 6 (Right). For a wide range of parameter values the top-ranked solver (“SBVA-sbva_cadical”) and the 4th-ranked solver (“Kissat_MAB_prop-no.sym”) are best according to the exponential utility function as well.

5 Conclusion

We have proposed four improvements to the utilitarian algorithm configuration procedure COUP that boost its empirical performance without degrading its theoretical bounds. We have shown that these improvements make a real difference to the tightness of the bounds proved by COUP, as well as to the quality of the returned configuration. We demonstrated that COUP is competitive with SMAC, the most widely used heuristic configuration procedure, while also offering theoretical optimality guarantees that improve with time. We are now inclined to believe that utilitarian algorithm configuration methods that offer theoretical guarantees have reached the point that they should be preferred for use in practice. We also described a case study on the SAT Competition that shows how to reason about the robustness of a conclusion recommended by a utilitarian configuration procedure with respect to different utility functions that could have been chosen, implying that uncertainty about the exact utility function used need not be a barrier to the adoption of utilitarian algorithm configuration approaches. In future work we plan to perform a more extensive experimental evaluation to make this claim definitive, comparing COUP’s performance with SMAC across a wider set of utility functions, and also contrasting COUP with other heuristic methods.

Acknowledgments

We would like to thank Greg d'Eon for helpful discussions and suggestions regarding plots. 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

- Ansótegui, C.; Sellmann, M.; and Tierney, K. 2009. A gender-based genetic algorithm for the automatic configuration of algorithms. In *International Conference on Principles and Practice of Constraint Programming*, 142–157. Springer.
- Birattari, M.; Stützle, T.; Paquete, L.; and Varrentrapp, K. 2002. A Racing Algorithm for Configuring Metaheuristics. In *Proceedings of the 4th Annual Conference on Genetic and Evolutionary Computation*, 11–18. Morgan Kaufmann Publishers Inc.
- Brandt, J.; Schede, E.; Haddenhorst, B.; Bengs, V.; Hüllermeier, E.; and Tierney, K. 2023. Ac-band: A combinatorial bandit-based approach to algorithm configuration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 12355–12363.
- Cappé, O.; Garivier, A.; Maillard, O.-A.; Munos, R.; and Stoltz, G. 2013. Kullback-Leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, 1516–1541.
- Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.
- Gabillon, V.; Ghavamzadeh, M.; and Lazaric, A. 2012. Best arm identification: A unified approach to fixed budget and fixed confidence. *Advances in neural information processing systems*, 25.
- Graham, D.; and Leyton-Brown, K. 2024. Utilitarian Algorithm Configuration for Infinite Parameter Spaces. *arXiv preprint arXiv:2405.18246*.
- Graham, D. R.; Leyton-Brown, K.; and Roughgarden, T. 2023a. Formalizing preferences over runtime distributions. In *International Conference on Machine Learning*, 11659–11682. PMLR.
- Graham, D. R.; Leyton-Brown, K.; and Roughgarden, T. 2023b. Utilitarian Algorithm Configuration. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Hoeffding, W. 1963. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 53(301): 13–30.
- Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2011. Sequential model-based optimization for general algorithm configuration. In *International conference on learning and intelligent optimization*, 507–523. Springer.
- Hutter, F.; Hoos, H. H.; Leyton-Brown, K.; and Stützle, T. 2009. ParamILS: an automatic algorithm configuration framework. *Journal of Artificial Intelligence Research*, 36: 267–306.
- Hutter, F.; López-Ibáñez, M.; Fawcett, C.; Lindauer, M.; Hoos, H. H.; Leyton-Brown, K.; and Stützle, T. 2014. Aclib: A benchmark library for algorithm configuration. In *International Conference on Learning and Intelligent Optimization*, 36–40. Springer.
- Jamieson, K.; and Talwalkar, A. 2016. Non-stochastic best arm identification and hyperparameter optimization. In *Artificial intelligence and statistics*, 240–248. PMLR.
- Kalyanakrishnan, S.; Tewari, A.; Auer, P.; and Stone, P. 2012. PAC subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, 655–662.
- Kleinberg, R.; Leyton-Brown, K.; and Lucier, B. 2017. Efficiency Through Procrastination: Approximately Optimal Algorithm Configuration with Runtime Guarantees. In *IJCAI*, volume 3, 1.
- 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. L. 1987. Adaptive treatment allocation and the multi-armed bandit problem. *The annals of statistics*, 1091–1114.
- Lai, T. L.; 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.
- Li, L.; Jamieson, K.; DeSalvo, G.; Rostamizadeh, A.; and Talwalkar, A. 2018. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, 18(185): 1–52.
- Lindauer, M.; Eggenberger, K.; Feurer, M.; Biedenkapp, A.; Deng, D.; Benjamins, C.; Ruhkopf, T.; Sass, R.; and Hutter, F. 2022. SMAC3: A Versatile Bayesian Optimization Package for Hyperparameter Optimization. *Journal of Machine Learning Research*, 23: 1–9.
- López-Ibáñez, M.; Dubois-Lacoste, J.; Cáceres, L. P.; Birattari, M.; and Stützle, T. 2016. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3: 43–58.
- Roussel, O. 2011. Controlling a Solver Execution with the runsolver Tool: System description. *Journal on Satisfiability, Boolean Modelling and Computation*, 7(4): 139–144.
- Weisz, G.; György, A.; Lin, W.-I.; Graham, D.; Leyton-Brown, K.; Szepesvári, C.; and Lucier, B. 2020. Impatient-CapsAndRuns: Approximately optimal algorithm configuration from an infinite pool. *Advances in Neural Information Processing Systems*, 33: 17478–17488.
- Weisz, G.; György, A.; and Szepesvári, C. 2018. LeapsAndBounds: A method for approximately optimal algorithm configuration. In *International Conference on Machine Learning*, 5257–5265. PMLR.

Weisz, G.; Gyorgy, A.; and Szepesvári, C. 2019. CapsAndRuns: An improved method for approximately optimal algorithm configuration. In *International Conference on Machine Learning*, 6707–6715. PMLR.

Xu, L.; Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2008. SATzilla: portfolio-based algorithm selection for SAT. *Journal of artificial intelligence research*, 32: 565–606.