

## Using machine learning to predict algorithm runtime.

BY KEVIN LEYTON-BROWN, HOLGER H. HOOS, FRANK HUTTER, AND LIN XU

# Understanding the Empirical Hardness of $NP$ -Complete Problems

PROBLEMS ARE INTRACTABLE when they “can be solved, but not fast enough for the solution to be usable.”<sup>13</sup>  $NP$ -complete problems are commonly said to be intractable; however, the reality is more complex. All known algorithms for solving  $NP$ -complete problems require exponential time in the worst case; however, these algorithms nevertheless solve many problems of practical importance astoundingly quickly, and are hence relied upon in a broad range of applications. The propositional satisfiability problem (SAT) serves as a good example. One of the most popular approaches for the formal verification of hardware and software relies on general-purpose SAT solvers and SAT encodings, typically with hundreds of thousands of variables. These instances can often be solved in

seconds, even though the same solvers can be stymied by handcrafted instances involving only hundreds of variables.

Clearly, we could benefit from a more nuanced understanding of algorithm behavior than is offered by asymptotic, worst-case analysis. Our work asks the question most relevant to an end user: “*How hard is it to solve a given family of problem instances, using the best available methods?*” Formal, complexity-theoretic analysis of this question seems hopeless: the best available algorithms are highly complex (and, in some cases, only available in compiled form), and instance distributions representative of practical applications are heterogeneous and richly structured. For this reason, we turn to statistical, rather than combinatorial, analysis.

The main claim of this article is that rigorous statistical methods can characterize algorithm runtime with high levels of confidence. More specifically, this article surveys over a decade of research showing how to build *empirical hardness models (EHMs)* that, given a new problem instance, estimate the runtime of an algorithm in low-order polynomial time.<sup>14,16,18,21,26–29,32,33,39,40</sup> We have shown that it is possible to build quite accurate models for different  $NP$ -complete problems (we have studied SAT, combinatorial auction winner determination, mixed integer programming, and the traveling salesman problem), distributions of prob-

### » key insights

- Rigorous statistical methods can characterize algorithm runtime with high levels of confidence.
- The resulting models can be analyzed to gain insight into what makes problems hard, even in settings where theoretical analysis cannot get traction.
- Accurately predicting runtime has many practical applications, such as selecting the best candidate from an algorithm portfolio, making benchmark distributions harder, and finding parameter settings that optimize an algorithm's performance.


lem instances (we have considered dozens), and solvers (again, dozens). We have robustly found that even very succinct EHMs can achieve high accuracies, meaning they describe simple relationships between instance characteristics and algorithm runtime.<sup>a</sup> This makes our approach important even for theoretically inclined computer scientists who prefer proofs to experimental findings: EHMs can uncover new, simple relationships between instance characteristics and runtime, and thereby catalyze new theoretical work.

The focus of this article is on ways that EHMs contribute to our *understanding* of *NP*-complete problems; however, they are also useful in a variety of practical applications. Most straightforwardly, they can aid the distribution of problem instances across a cluster, or predict how long a run will take to complete. More interestingly, they can be used to combine a set of high-variance algorithms into an “algorithm portfolio” that outperforms its constituents; be leveraged to automatically make benchmark distributions more challenging; and aid in the configuration (or “tuning”) of highly parameterized algorithms for good performance on given instance distributions. More detailed explanations of these applications appear in sidebars throughout this article.

### Phase Transitions in Uniform-Random 3-SAT

We begin by describing the most widely known relationship between a characteristic of fixed-size random SAT instances and solver runtime. (After this, we consider more realistic instances of SAT and other *NP*-hard problems.) Let  $p(c, v)$  denote the probability that a satisfiable 3-SAT formula<sup>b</sup>

<sup>a</sup> We do not survey the literature on algorithm performance prediction here; instead we focus on our own work. For extensive discussions of related work, please see Hutter et al.<sup>21</sup> and Leyton-Brown et al.<sup>29</sup>



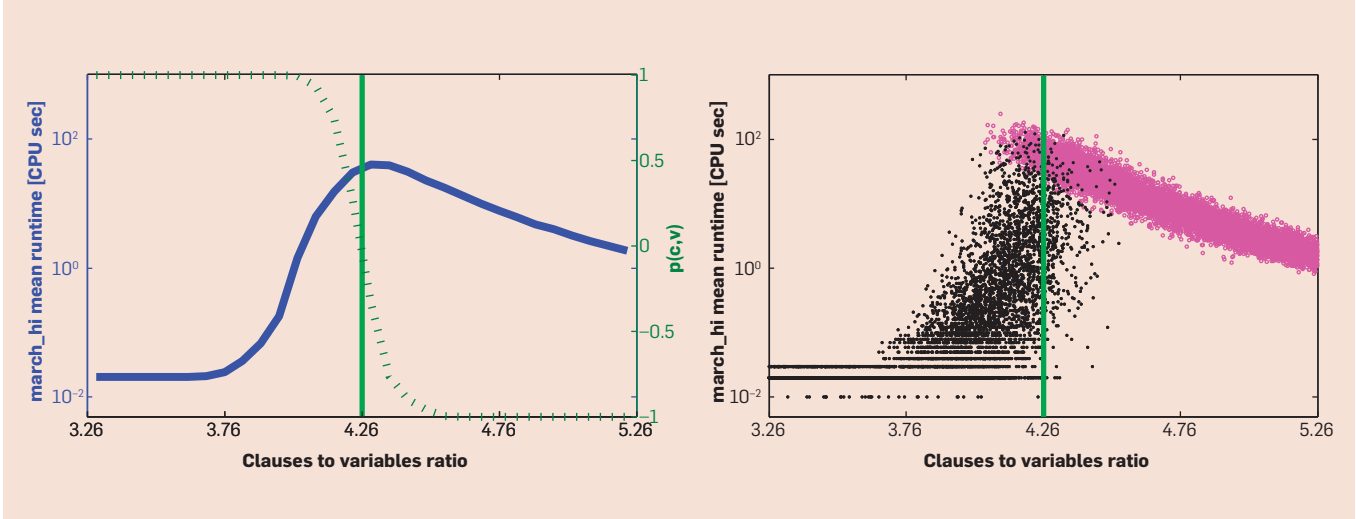
will be generated by uniformly sampling  $c$  clauses of three variables each from a set of  $v$  variables, negating each with probability 0.5. In the early 1990s, researchers discovered that when  $v$  is held constant,  $p(c, v)$  exhibits a “phase transition” as  $c/v$  crosses a critical value of about 4.26.<sup>8,31</sup> Intuitively, instances with few clauses are underconstrained and thus almost always satisfiable, while those with many clauses are overconstrained and almost always unsatisfiable. The interesting fact is that, for all fixed values of  $v$  so far tested, the phase transition

<sup>b</sup> A SAT formula  $F$  is solved by deciding whether there exists an assignment of its variables under which  $F$  evaluates to true. A subclass of particular importance is 3-SAT. A 3-SAT instance is a conjunction of *clauses*, each of which is a disjunction of three variables or their negations. For example,  $(v_1 \vee \neg v_2 \vee v_4) \wedge (\neg v_1 \vee \neg v_3 \vee v_4)$  is a simple formula with  $v = 4$  variables and  $c = 2$  clauses that has several satisfying assignments (for example,  $[v_1, v_2, v_3, v_4] = [\text{true}, \text{true}, \text{false}, \text{false}]$ ).

point at which  $p(c, v)$  is exactly 0.5, appears to coincide with a runtime peak even for the SAT solvers that perform best on these instances. This finding thus links an algorithm-independent property of an instance ( $c/v$ ) with algorithm-specific runtime in a way that has proven robust across solvers.

Figure 1 (left) shows this relationship using real data. The dotted line shows  $p(c, v)$  for uniform-random 3-SAT instances with  $v = 400$ , while the solid line shows the mean runtime of `march_hi`,<sup>11</sup> one of the best SAT solvers for uniform-random 3-SAT, on the same instances. We do indeed observe both a phase transition and a hardness spike at the phase transition point. However, there is more to the story. Figure 1 (right) plots raw runtime data (on a log scale) for `march_hi`, with each point corresponding to a single (random) 3-SAT formula. We can now see that the  $c/v$  ratio does not suffice to fully explain `march_hi`'s empirical behavior on these instances: there is still substan-

**Figure 1. Runtime of `march_hi` on uniform-random 3-SAT instances with  $v = 400$  and variable  $c/v$  ratio. Left: mean runtime, along with  $p(c, v)$ ; right: per-instance runtimes, colored by satisfiability status. Runtimes were measured with an accuracy of 0.01s, leading to the discretization effects visible near the bottom of the figure. Every point represents one SAT instance.**



tial variation at each point along the  $x$  axis—over two orders of magnitude at the “hard” phase transition point. The runtime pattern also depends on satisfiability status: hard instances are scarcer and runtime variation is greater among satisfiable instances than among unsatisfiable instances. One reason for this is that on satisfiable instances the solver can stop as soon as it encounters a satisfying assignment, whereas for unsatisfiable instances a solver must prove that no satisfying assignment exists anywhere in the search tree.

### A Case Study on Uniform-Random 3-SAT

We now ask whether we can better understand the relationship between instance structure and solver runtime by considering instance features beyond just  $c/v$ . We will then use a machine learning technique to infer a relationship between these features and runtime. Formally, we start with a set  $I$  of instances, a vector  $x_i$  of feature values for each  $i \in I$ , and a runtime observation  $y_i$  for each  $i \in I$ , obtained by running a given algorithm on  $i$ . Our goal will be to identify a mapping  $f: x \rightarrow y$  that predicts  $y_i$  as accurately as possible, given  $x_i$ . We call such a mapping an EHM.<sup>c</sup> Observe that we have just

described a supervised learning problem, and more specifically a regression problem. There are many different regression algorithms that one could use to solve this problem, and indeed, over the years we have considered about a dozen alternatives. Later in this article we will advocate for a relatively sophisticated learning paradigm (random forests of regression trees), but we begin by discussing a very simple approach: quadratic ridge regression.<sup>5</sup> This method performs linear regression based on the given features and their pairwise products, and penalizes increases in feature coefficients (the “ridge”). We elaborate this method in two ways. First, we transform the response variable by taking its (base-10) logarithm; this better allows runtimes, which vary by orders of magnitude, to be described by a linear model. Second, we reduce the set of features by performing forward selection: we start with an empty set and iteratively add the feature that (myopically) most improves prediction. The result is simpler, more robust models that are less prone to numerical problems. Overall, we have found that even simple learning algorithms like this one usually suffice to build strong EHMs; more important is identifying a good set of instance features.

**Instance features.** It can be difficult to identify features that correlate as strongly with instance hardness as  $c/v$ . We therefore advocate including all features that show some promise of being predictive, and relying on the machine-

learning algorithm to identify the most useful ones. Our only requirement is that the features be computable in low-order polynomial time; in some applications, we also restrict ourselves to features that are quadratic time or faster. For the SAT domain, we defined 138 features summarized as:

- ▶ Problem size measures  $c$  and  $v$ , plus nonlinear combinations we expected to be important, like  $c/v$  and  $c/v - 4.26$ ;
- ▶ Syntactic properties of the instance (proximity to Horn clauses, balance of positive and negative literals, and so on);
- ▶ Constraint graph statistics. We considered three graphs: nodes for variables and edges representing shared constraints (clauses); nodes for clauses and edges representing shared variables with opposite polarity; nodes for both clauses and variables, and edges representing the occurrence of a variable in a given clause. For each graph, we computed various statistics based on node degrees, path lengths and clustering, among others;
- ▶ A measure of the integrality of the optimal solution to the linear programming relaxation of the given SAT instance—specifically, the distance between this solution and the nearest (feasible or infeasible) integral point;
- ▶ Knuth’s estimate of search tree size;<sup>25</sup> and,
- ▶ Probing features computed by running bounded-length trajectories of local search and tree search algorithms and extracting statistics from

<sup>c</sup> It is sometimes useful to build EHMs that predict a probability distribution over runtimes rather than a single runtime; see Hutter et al.<sup>21</sup> For simplicity, here we discuss only the prediction of mean runtime.

these probes (for example, number of steps before reaching a local minimum in local search or amount of unit propagation performed in tree search).

### Model Performance

Let us now investigate the models we can build using these techniques for uniform-random 3-SAT. We consider two sets of instances: one in which the  $c/v$  ratio varies around the phase transition point, and another in which it is fixed at  $c/v = 4.26$ . The first set of instances is less challenging, as we already know the  $c/v$  ratio suffices to explain much of the runtime variation. However, this set is still useful as a sanity check to ensure our methods discover the importance of the  $c/v$  feature, and to investigate what additional features turn out to be useful. The second set contains fixed-size instances in the hard  $c/v = 4.26$  region: any patterns we can find here are interesting since we cannot distinguish instances based on their  $c/v$  ratio. In both cases (as with all other empirical results we have shown) we randomly partitioned our data into a “training set” we used to build the EHM and a disjoint “test set” we used solely to evaluate the performance of the EHM, thereby assessing the accuracy of a model’s predictions beyond the data used for constructing it.

Figure 2 shows the results of our investigation, giving true versus predicted `march_hi` runtimes, with every point in the figure corresponding to a different test-set problem instance. Overall, the points cluster around the

diagonal in both plots, meaning the predictions are reasonably accurate. (Indeed, these results are better than they appear to the eye, as a greater density of points falls closer to the diagonal.) Observe that prediction accuracy was considerably better for unsatisfiable instances than for satisfiable instances. Root mean squared error (RMSE) is a numerical measure of a model’s accuracy; the two models achieved RMSEs of 0.31 and 0.56, respectively. A model that consistently mispredicted runtimes by a factor of 10 would have achieved an RMSE of 1.0.

We used the variable ratio instance set to verify that the model recognizes the importance of the  $c/v$  feature; similarly, we would like to identify other informative features. We cannot do this simply by examining a scatterplot, nor by looking at the coefficients of the model itself: because many of our features are strongly correlated, important features can have small coefficients and unimportant features can have big coefficients. Instead, we identify new models that use only a small number of uncorrelated features by applying the same forward selection method described previously, but terminating much earlier, when the benefit from adding new features begins to taper off. Then (following Friedman<sup>10</sup>) we quantify the importance of each feature in these new models by measuring the loss in RMSE incurred by omitting it, and scaling these values to give the most important feature receives a score of 100.

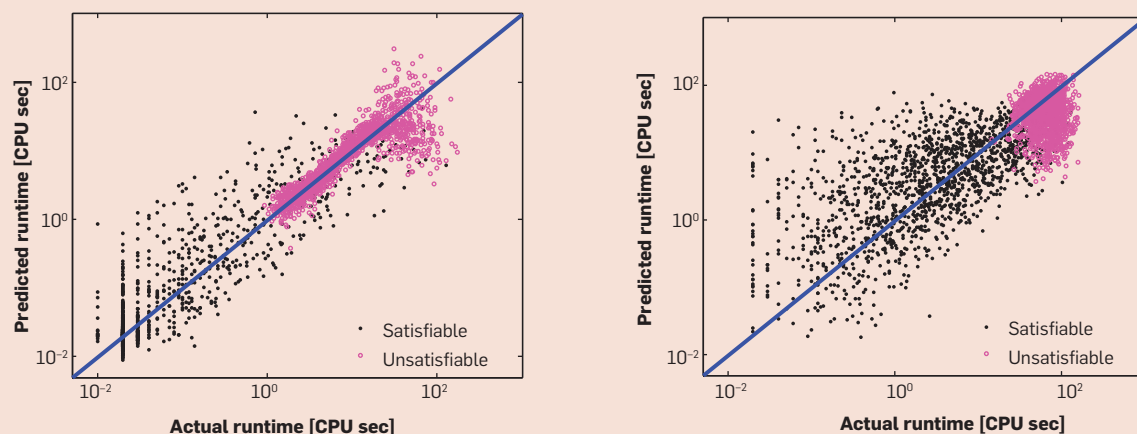
Table 1 demonstrates that the model indeed identified  $c/v$  (and the variant  $|c/v - 4.26|$ ) as an important feature; recall that our quadratic regression has access to all pairwise products of features. Other features were also important, in particular, `SAPS_BestSolution_Mean`, the average number of satisfied clauses achieved by short runs of the local search procedure SAPS.<sup>20</sup> This is interesting, as one might not have expected the performance of a local search algorithm to be informative about the performance of a tree search algorithm.

For the fixed-ratio data,  $c/v$  was constant, and we therefore see a different picture among important features (Table 2). Again, local search probing features figured prominently (`GSAT_BestSolution_Mean` gives the average number of satisfied clauses achieved by short runs of `GSAT`<sup>36</sup>). Another important feature was `CG_entropy`. It gives the entropy across node degrees in the “constraint graph” in which nodes correspond to clauses and edges indicate that a pair

**Table 1. Feature importance,  $c/v \in [3.26, 5.26]$ .**

Feature	Score
$ c/v - 4.26  \times \text{SAPS\_BestSolution\_Mean}$	100
$c/v \times \text{SAPS\_BestSolution\_Mean}$	19
$\text{GSAT\_BestSolution\_CoeffVar} \times \text{SAPS\_BestStep\_CoeffVar}$	19
$\text{SAPS\_BestStep\_CoeffVar} \times \text{CG\_entropy}$	18

**Figure 2. Actual vs. predicted runtimes for `march_hi` on uniform-random 3-SAT. Each dot represents a test instance not used to train the model; perfect predictions would fall along the diagonal. Left:  $c/v \in [3.26, 5.26]$ ; Right:  $c/v = 4.26$ .**



of clauses share one or more variables of opposite sign. We have repeated this analysis for other SAT solvers (for example, `kcnfs` and `sat3`) and obtained the same qualitative results: runtime is predictable with high accuracy, small models suffice for good performance, and local-search and constraint-graph features are important.<sup>33</sup>

We have already observed that satisfiable and unsatisfiable instances exhibit very different distributions of algorithm runtimes. We thus considered the problem of building EHMs only for satisfiable or for unsatisfiable instances. (We call these “conditional models” as they depend on knowing the satisfiability of a given instance.) We found that conditional EHMs were more accurate than unconditional EHMs; more interestingly, *single-feature* conditional models turned out to be sufficient for predicting runtime with high accuracy. For satisfiable instances, this feature was `GSAT_BestSolution_Mean`, whereas for unsatisfiable instances, it was `Knuth_Mean`. We can explain these findings as follows. Since local

search algorithms apply heuristics to find a solution as quickly as possible, the reliability with which such an algorithm is able to make quick progress is informative about the speed at which a complete algorithm will be able to find a solution. Tree search algorithms must rule out every node in the tree to prove unsatisfiability; thus, an estimate of this tree size is the most important feature in the unsatisfiable case.

**Predicting Satisfiability Status**

These observations led us to a new idea: building a classifier that directly predicts satisfiability status, and then leveraging conditional EHMs based on this prediction. There are two reasons to be skeptical about this approach. First, conditional EHMs can make very inaccurate predictions on instances with the “wrong” satisfiability status, so it is not clear that we would obtain improved accuracy overall. Second, and more fundamentally, it may seem doubtful that we could accurately predict satisfiability status—that would correspond to guessing whether an instance is solvable without actually solving it! Despite this reservation, and applying sophisticated statistical techniques to mitigate the potential cost of mispredictions (see Xu et al.<sup>38</sup>), we did build *hierarchical hardness models* for uniform random 3-SAT instances. These models achieved (relatively modest) improvements in predictive accuracy on both the fixed-ratio and variable-

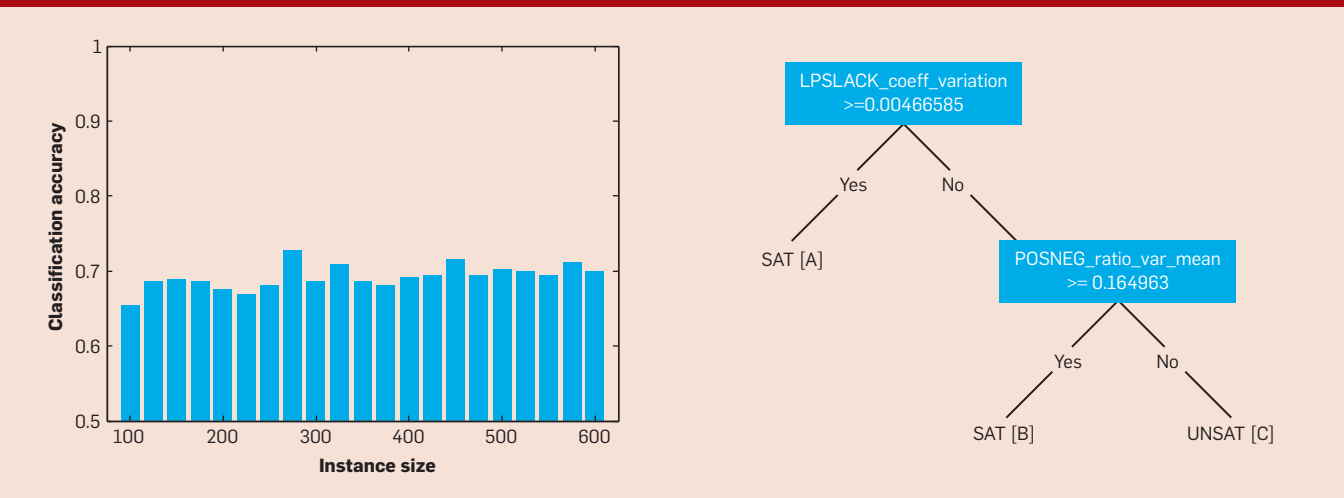
ratio instance sets. Clearly, hierarchical hardness models can only outperform regular EHMs if our classifier is able to accurately predict satisfiability status. In the variable-ratio case, a natural baseline is a classifier that predicts satisfiability for instances with  $c/v < 4.26$ , and unsatisfiability otherwise. This classifier achieved an accuracy of 96%, underlining the predictive power of the  $c/v$  feature. Our classifier further increased this accuracy to 98% (halving the error rate). Unlike the baseline, our classifier also applies in the fixed-ratio case, where it achieved accuracy of 86%.

We found this result surprising enough that we investigated it in more depth.<sup>39</sup> We considered instances of varying size, from 100 variables (solvable in milliseconds) to 600 variables (solvable in a day; about the largest instances we could solve practically). We focused on instances generated at the phase transition point, because they pose the hardest prediction problem: the probability of generating satisfiable instances at this point is 50%, and in practice, our data sets were indeed very evenly balanced between satisfiable and unsatisfiable instances. Our aims were to investigate whether prediction accuracy appeared to fall to that of random guessing on larger problems, and if not, to give an easily comprehensible model that could serve as a starting point for theoretical analysis. In doing so we restricted our models in three ways that each reduced predictive accuracy, but allowed

**Table 2. Feature importance,  $c/v = 4.26$ .**

Feature	Score
<code>GSAT_BestSolution_Mean<sup>2</sup></code>	100
<code>GSAT_BestSolution_Mean</code>	88
<code>SAPS_BestSolution_CoeffVar</code> <code>× SAPS_AvgImprove_Mean</code>	33
<code>SAPS_BestStep_CoeffVar</code> × <code>CG_Entropy</code>	22

**Figure 3. Left: Classification accuracies for our simple decision tree on uniform-random 3-SAT instances at the phase transition with varying numbers of variables. The tree was trained only on 100-variable data. Right: The decision tree. Predictive accuracies for instances falling into the three regions were between 60% and 70% (region A); about 50% (region B); and between 70% and 80% (region C).**



us to better answer these questions. First, we permitted ourselves to train our classifier only on the 100-variable instances. Second, we considered only decision trees<sup>7</sup> having at most two decision nodes. (We obtained these models by a standard decision tree learning procedure: greedily choosing a feature that best partitioned the training data into satisfiable and unsatisfiable instances, and recursively splitting the resulting partitions.) Finally, we omitted all probing features. Although probing (for example, local-search) features were very useful for prediction, they were disproportionately effective on small instances and hence would have complicated our study of scaling behavior. Also, because we aimed to obtain easily comprehensible models, we were disinclined to use features based on complex, heuristic algorithms.

Given all of these restrictions, we were astonished to observe predictive accuracies consistently above 65%, and apparently independent of problem size (see Figure 3, left; statistical testing showed no evidence that accuracy falls with problem size). Indeed, our first two restrictions appear to have come at low cost, decreasing accuracies by only about 5%. (Furthermore, after lifting these restrictions, we still found no evidence that accuracy was affected by problem size.) Hence, the reader may be interested in understanding our two-feature model in more detail; we hope it will serve as a starting point for new theoretical analysis of finite-size SAT instances at the phase transition. The model is given in Figure 3 (Right). `LPSLACK_coeff_variation` is based on solving a linear programming relaxation of an integer program representation of SAT instances. For each variable  $i$  with LP solution value  $S_i \in [0, 1]$ , `LPSLACKi` is defined as  $\min\{1 - S_i, S_i\}$ : the deviation of  $S_i$  from integrality. `LPSLACK_coeff_variation` is then the coefficient of variation (the standard deviation divided by the mean) of the vector `LPSLACK`. The  $i$ th element of the vector `POSNEG_ratio_var_mean` is the average ratio of positive and negative occurrences of each variable. For each variable  $i$  with  $P_i$  positive occurrences and  $N_i$  negative occurrences, `POSNEG_ratio_var` is defined as  $|0.5 - P_i / (P_i + N_i)|$ . `POSNEG_ratio_var_mean` is then

the average over elements of the vector `POSNEG_ratio_var`. Finally, recall that our model was trained on constant-size instances; we normalized the `LPSLACK_coeff_variation` and `POSNEG_ratio_var_mean` features to have mean 0 and standard deviation 1 on this training set. To evaluate the model on a given instance of a different size, we randomly sampled many new instances of that size to compute new normalization factors, which we then applied to the given instance.

### Beyond Uniform-Random 3-SAT

Our original motivation involved studying real problems faced by a practitioner, problems that are very unlikely to have uniform random structure. Thus, it is important to demonstrate that EHMs work reliably for a wide range of more realistic instance distributions, and that they are not limited to SAT. In short, they do, and they are not. By now, we have built EHMs for four different NP-complete problems: SAT<sup>14,16,21,33,38,40</sup> the combinatorial auction winner determination problem (WDP),<sup>28,29</sup> mixed integer programming (MIP, a standard encoding for problems with both discrete and

continuous variables),<sup>14,19,21</sup> and the traveling salesman problem (TSP).<sup>21</sup> Observe that we have considered both optimization and decision problems, and that these problems involve discrete variables, continuous variables, and combinations of the two. For each problem, we derived a new set of instance features. This was not trivial, but not terribly difficult either; in all cases we used problem size measures, syntactic properties, and probing features. Extending what we know now to a new domain would probably entail a few days of work. In our publications and other technical work (for example, submissions to SAT competitions) we have considered more than 30 instance distributions. These include sophisticated random generators (for example, SAT reductions from graph coloring and factoring; combinatorial auction benchmarks based on economic models); instance sets from public benchmarks and competitions (for example, MIPLIB; the SAT competition); and sets of instances derived from practical applications (for example, SAT-encoded instances from software verification and bounded model checking; industrial MIP instances ranging from

## Application 1. Algorithm Selection (SATzilla)

There currently exists no “best” SAT solver; different solvers perform well on different families of instances, and performance differences between them are typically very large. The effectiveness of EHMs suggests a straightforward solution to the *algorithm selection problem*:<sup>35</sup> given a new problem instance, predict the runtime of several SAT solvers, and then run the one predicted to be fastest. This approach<sup>27</sup> forms the core of SATzilla<sup>32,33,40</sup> a portfolio-based algorithm selector for SAT.

SATzilla first participated in the 2003 SAT Competition (<http://www.satcompetition.org>), and placed second and third in several categories. We have since extensively improved the approach, allowing randomized and local search algorithms as component solvers; introducing the idea of presolvers that are run for a short, fixed time before the selected solver; adding the ability to optimize for complex scoring functions; and automating the construction of the selector (for example, pre-solver selection; component solver selection) given data. Leveraging these improvements, and benefiting from the continued improvement of the component solvers upon which it draws, SATzilla led the field in the 2007 and 2009 SAT Competitions, winning five medals each time.

More recently, our design of SATzilla evolved from selection based on runtime predictions (EHMs) to a cost-sensitive classification approach that directly selects the best-performing solver without predicting runtime.<sup>41</sup> In the 2012 SAT Challenge (<http://baldur.iti.kit.edu/SAT-Challenge-2012>), SATzilla was eligible to enter four categories; it placed first in three of these and second in the fourth. Overall, SATzilla’s success demonstrates the effectiveness of automated, statistical methods for combining existing solvers—including “uncompetitive” solvers with poor average performance. Except for the instance features used by our models, our approach is entirely general, and is likely to work well for other problems with high runtime variation. All of our software is publicly available; see <http://www.cs.ubc.ca/labs/beta/Projects/SATzilla>.

machine job allocation to wildlife conservation planning; TSP instances representing drilling circuit boards and traveling between real cities). We have also studied more than 50 state-of-the-art solvers, both open source projects and proprietary tools developed by industry. Our solvers were both deterministic and randomized, and both complete (that is, guaranteed to find a solution if one exists) and incomplete. In many cases, we only had access to an executable of the solver, and in no case

did we make use of knowledge about a solver's inner workings. As mentioned earlier, we have also gone beyond quadratic basis function regression to study more than a dozen other statistical modeling techniques, including lasso regression, multivariate adaptive regression splines, support vector machine regression, neural networks, Gaussian processes, regression trees and random forests (see Leyton-Brown et al.<sup>29</sup> and Hutter et al.<sup>21</sup>). We omit the details here, but state the conclusion:

we now prefer random forests of regression trees,<sup>6</sup> particularly when the instance distribution is heterogeneous. We briefly describe this model class for completeness, but refer readers to the literature for details.<sup>6,7</sup> Regression trees are very similar to decision trees (which we used here for predicting satisfiability status). However, as a classification method, decision trees associate categorical labels with each leaf (for example, “satisfiable,” “unsatisfiable”), while regression trees associate a real-valued prediction with each leaf. Random forests report an average over the predictions made by each of an ensemble of regression trees; these trees are made to differ by randomizing the training process.

Figure 4 illustrates the results of our broader experience with EHMs by highlighting three different solvers, each from a different domain. In each case we give plots for both quadratic ridge regression and random forests to show the impact of the learning algorithm. First (column 1), we considered the prominent SAT solver, Minisat 2.09, running on a very heterogeneous mix of instances from the international SAT competition. Whereas the competition subdivides instances into categories (“industrial/application,” “handmade/crafted,” and “random”), we merged all instances together. Likely because of the heterogeneity of the resulting set, quadratic regression performed relatively poorly here. Random forests yielded much more reliable estimates; notably, they can partition the feature space into qualitatively different parts, and they never predict runtimes larger or smaller than the extrema observed in the training data. However, observe that even the less accurate quadratic regression models were usually accurate enough to differentiate between fast and slow runs in this domain; see the sidebar on SATzilla. Second (column 2), we studied the performance of the leading complete TSP solver, Concorde,<sup>2</sup> on a widely used suite of rather homogeneous, randomly generated TSP instances.<sup>23</sup> We again see good performance, now for both quadratic regression and random forests. Third (columns 3 and 4), to show the effect of changing only the instance distribution, we consider one solver on two different distributions. IBM ILOG

## Application 2. Generating Hard Benchmarks

Realistic, hard benchmark distributions are important because they are used as an objective measure of success in algorithm development. However, it can sometimes be just as difficult to find new, hard benchmarks as it is to find new strategies for solving previously hard benchmarks. To fill this gap, EHMs can be used to automatically adjust existing instance generators so they produce instances that are harder for a given set of algorithms.<sup>26,29</sup>

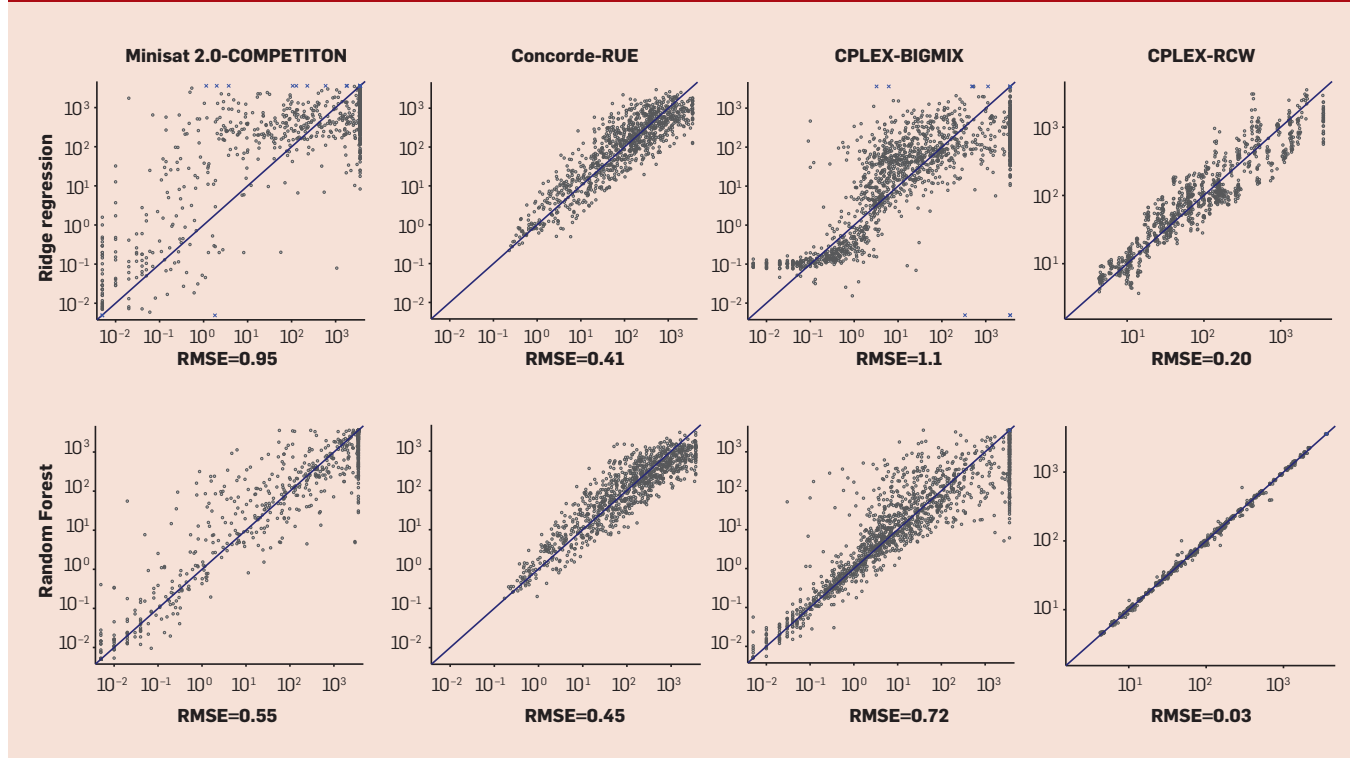
We start with an instance generator that has parameters  $p$ . Such generators are often simply used with default parameter settings; however, to search for harder instances we instead sample each parameter's value uniformly from a fixed range. Call the resulting distribution over instances  $\mathcal{D}$ . Our goal is to sample from a new distribution  $\mathcal{D}'$  over the same instances that weights instances by their hardness for a given algorithm  $A$ . (Think of  $A$  as having the best average runtime among all algorithms in a given set, or as being a SATzilla-style selector among such algorithms.) We can do this via a form of importance sampling. We construct an EHM for  $A$  on  $\mathcal{D}$  using our standard instance features  $f$ ; for an instance  $x$ , call this model's prediction  $H_f(x)$ . We would like to generate a large set of instances from  $\mathcal{D}$ , weight each instance  $x$  in proportion to  $H_f(x)$ , and then sample a single instance from the set in proportion to the weights. This approach works, but requires a very large number of samples when hard instances are rare in  $\mathcal{D}$ . To improve performance, we learn a quadratic EHM  $H_p$  that uses only the generator parameters  $p$  as features. We can then sample instances  $x$  in proportion to  $\mathcal{D}(x) \cdot H_p(x)$  rather than sampling from  $\mathcal{D}$  (by sampling directly from polynomial function  $H_p$ , and then running the instance generator with the resulting parameters), and then weight each sampled instance  $x$  by  $H_f(x) / H_p(x)$ .  $H_p$  thus guides our search towards harder instances without biasing the weights. In experiments with the Combinatorial Auction Test Suite<sup>30</sup> this approach increased the mean hardness of generated instances by up to a factor of 100,<sup>26,29</sup> and often created instances much harder than we had ever observed using the generators' default parameters.

## Application 3. Algorithm Configuration

Imagine that each time the designer of a heuristic algorithm faced a choice about a given design element, she simply encoded it as a free parameter of a single solver. In the end, her problem of designing a good algorithm for a given problem domain would be reduced to the stochastic optimization problem of finding a configuration that achieved good performance.<sup>12,14,24</sup>

We have applied automated methods for solving this problem to identify novel algorithm configurations that have yielded orders of magnitude speedups in a broad range of domains, including SAT-based formal verification,<sup>15</sup> MIP solving,<sup>17</sup> and automated planning.<sup>37</sup> One state-of-the-art method for solving this problem is based on EHMs: *sequential model-based algorithm configuration* (SMAC)<sup>19</sup> iterates between using an EHM to select promising configurations to explore next, executing the algorithm with these configurations, and updating the model with the resulting information. SMAC is freely available; please see <http://aclib.net/SMAC>. EHMs can also be used to select configurations on a per-instance basis.<sup>16</sup>

**Figure 4. Visual comparison of models for runtime predictions on unseen instances. In each plot, the x-axis denotes true runtime and the y-axis runtime as predicted by the respective model. Predictions above 3,000 or below 0.001 are denoted by a blue x.**



CPLEX<sup>22</sup> is the most widely used commercial MIP solver. BIGMIX is a highly heterogeneous mix of publicly available mixed integer programming problems. As with the mix of SAT instances in our first benchmark, linear regression struggled with predictions for some types of instances and occasionally made catastrophic mispredictions. Again, random forests performed much more robustly. RCW models the dispersal and territory establishment of the redcockaded woodpecker conditional on decisions about which parcels of land to protect.<sup>1</sup> The runtime of CPLEX for this domain was surprisingly predictable; random forests yielded among the best EHM performance we have ever observed.

### Beyond Single Algorithms

Unlike average-case complexity results that characterize the inherent complexity of a computational problem, EHMs always describe the performance of a given algorithm. In some sense this is an inherent limitation: a statistical approach cannot summarize the performance of algorithms that have not yet been invented. However, there is a useful way in which we can relax the single-algorithm restriction:

we can build a model that describes a *space* of existing algorithms.

More specifically, most state-of-the-art algorithms for hard combinatorial problems offer a range of algorithm parameters in order to enable users to customize or tune the algorithm's behavior. We define “parameters” very broadly, as encompassing any argument passed to a solver that changes its behavior (and, thus, its runtime) but not the nature of the solution it returns. Parameters can thus be continuous, categorical, ordinal, or Boolean, and can even be conditional on values taken by other parameters. More importantly, categorical and Boolean parameters can be used to represent very abstract decisions—effectively selecting among unrelated blocks of code—and can thereby open up vast algorithm design spaces. For example, IBM ILOG CPLEX exposes 76 parameters (45 categorical, six Boolean, 18 integer, and seven real-valued);<sup>17</sup> a fairly coarse discretization of these parameters yields over  $10^{47}$  different algorithm instantiations with vastly different performance profiles. We call such an instantiation of all parameters of a given algorithm to specific values a *configuration*. The second example

of a parameterized solver we use here is the SAT solver SPEAR,<sup>4</sup> which exposes 26 parameters (seven categorical, three Boolean, four integer, and 12 real-valued), giving rise to over  $10^{17}$  different algorithm instantiations.

We now consider generalizing EHMs to describe such parameterized algorithms. In principle, this is not much of a change: we consider models that map from a joint space of configurations and instance features to runtime predictions. The question is how well such an approach can work. Before we can give an answer, we need to decide how to evaluate our methods. We could test on the same configurations that we used to train the EHM but on new problem instances, on new configurations but on previously seen instances, or on combinations of previously unseen configurations and instances.

The third case is the hardest; it is the only setting for which we show results here. Figure 5 illustrates some representative results in this setting, focusing on random forest models. The first row shows scatterplots like those presented earlier, with each point representing a run of a randomly selected, previously unseen



configuration on a previously unseen instance. We also provide a different way of looking at the joint space of parameter configurations and instance feature vectors. The second row shows the true runtimes for each (configuration, instance) pair, sorting configurations by their average performance and sorting instances by their average hardness. Thus, the picture gives a snapshot of the runtime variation across both instances and configurations, and makes it possible to gauge how much of this variation is due only to differences between configurations versus differences between instances. Finally, the third row shows the predictions obtained from the EHM in the same format as the second row. This gives a way of visually comparing model performance to ground truth;

ideally, the second and third rows would look identical. (Indeed, when the two figures closely resemble each other, the EHM can serve as a *surrogate* for the original algorithm, meaning that the EHM can be substituted for the algorithm in an empirical analysis of the algorithm’s performance; see Hutter et al.<sup>18</sup>)

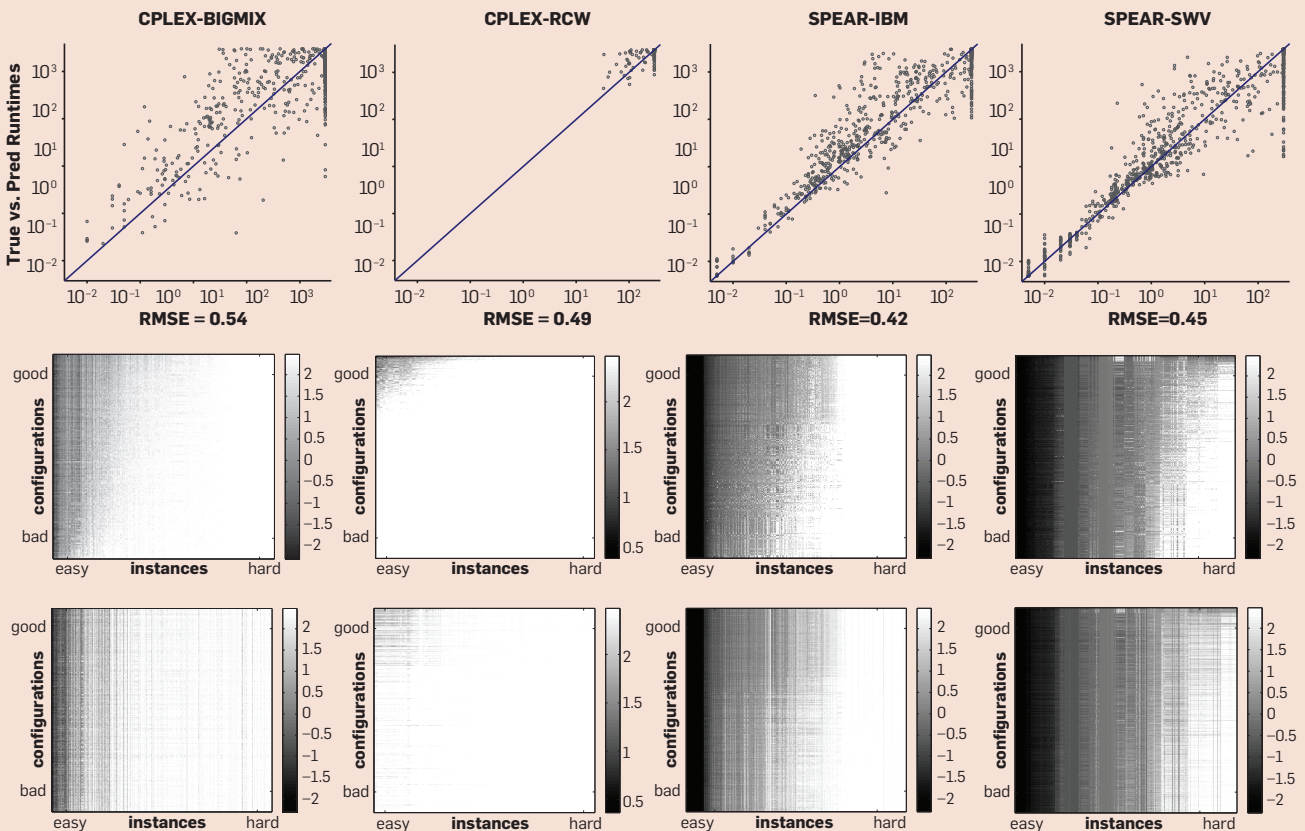
The main takeaway from our experiments is that our models were able to achieve high accuracies (RMSEs around 0.5; qualitative similarity between the second and third rows) even on algorithm configurations that were never examined during training. The first column of Figure 5 concerns runtime predictions for CPLEX on our heterogeneous mix of MIP instances. The instances in this benchmark differ greatly in hardness, much more

than the different CPLEX configurations differ in performance (see Row 2). As a result, instance hardness dominated runtimes and the model focused on the (more important) feature space, at the expense of failing to capture some of the performance difference between configurations. Second, on RCW most of the (randomly sampled) CPLEX configurations solved very few instances; we recorded such failures as very long runtimes. The model was nevertheless able to identify which configurations were good and which instances were easy. Finally, we consider predicting the runtime of SPEAR on two sets of formal verification instances: IBM is a set of bounded model checking instances,<sup>42</sup> while SWV is a set of software verification instances generated

**Figure 5. Visual comparison of models for runtime predictions on pairs of previously unseen test configurations and instances.**

Row 1: In each plot, the x-axis denotes true runtime and the y-axis denotes runtime as predicted by the respective model. Each dot represents one combination of an instance and parameter configuration, both previously unseen during model training.

Rows 2 and 3: Actual and predicted runtimes for each application domain. Each dot in the heat map represents the run of one parameter configuration on one instance; the grayscale value represents runtime on a log<sub>10</sub> scale (darker means faster).




with the Calysto static checker.<sup>3</sup> For both of these instance distributions, the runtime of SPEAR with different configurations was predicted with a high degree of accuracy. Our random forest models accurately predicted the empirical hardness of instances, the empirical performance of configurations, and even captured ways in which the two interact.

### Take-Away Messages

Statistical methods can characterize the difficulty of solving instances from a given distribution using the best available algorithms—even when those algorithms are extremely complex and traditional theoretical analysis is infeasible. Such EHMs are surprisingly effective in practice, across different hard combinatorial problems, real-world instance distributions, and state-of-the-art solvers. An analysis of these models can serve as a starting point for new theoretical investigations into complexity beyond the worst case, by identifying problem features that are predictive of hardness or that suffice to predict an objective function (for example, satisfiability status) directly. In the context of highly parameterized algorithms that span a large space of possible algorithm designs, we have found that it is even possible to predict the runtime of previously untested algorithm designs on previously unseen instances. EHMs have proven useful in a variety of practical applications, including the automatic design of algorithm portfolios, the automatic synthesis of hard benchmark distributions, and the automatic search for a performance-optimizing design in a large algorithm design space. We have written open source software for building EHMs, analyzing them, constructing algorithm portfolios, automatically configuring parameterized algorithms, and more: see <http://www.cs.ubc.ca/labs/beta/Projects/EPMS/>.

### Acknowledgments

Some work described in this article was performed with additional coauthors: Eugene Nudelman and Yoav Shoham made particularly sustained contributions, and Galen Andrew, Alex Devkar, and Jim McFadden also deserve mention. 

### References

- Ahmadizadeh, K., Dilkina, B., Gomes, C.P. and Sabharwal, A. An empirical study of optimization for maximizing diffusion in networks. In *Proceedings for Principles and Practice of Constraint Programming* (2010), 514–521.
- Applegate, D.L., Bixby, R.E., Chvátal, V. and Cook, W.J. *The Traveling Salesman Problem: A Computational Study*. Princeton University Press, 2006.
- Babić, D. and Hu, A.J. Structural abstraction of software verification conditions. In *Proceedings for Computer Aided Verification* (2007), 366–378.
- Babić, D. and Hutter, F. Spear theorem prover. Solver description. SAT 2007 Competition.
- Bishop, C.M. *Pattern Recognition and Machine Learning*. Springer, 2006.
- Breiman, L. Random forests. *Machine Learning* 45, 1 (2001), 5–32.
- Breiman, L., Friedman, J.H., Olshen, R. and Stone, C.J. *Classification and Regression Trees*. Wadsworth, Belmont, CA, 1984.
- Cheeseman, P., Kanefsky, B. and Taylor, W.M. Where the really hard problems are. In *Proceedings for International Joint Conference on Artificial Intelligence* (1991), 331–337.
- Eén, N. and Sörensson, N. An extensible SAT-solver. *Theory and Applications of Satisfiability Testing* (2004), 502–518.
- Friedman, J. Multivariate adaptive regression splines. *Annals of Statistics* 19, 1 (1991), 1–141.
- Heule, M. and Maaren, M.v. march\_hi. Solver description. SAT 2009 competition.
- Hoos, H.H. Programming by optimization. *Commun. ACM* 55, 2 (Feb. 2012), 70–80.
- Hopcroft, J.E., Motwani, R. and Ullman, J.D. *Introduction to Automata Theory, Languages, and Computation*. Pearson Education, 2007.
- Hutter, F. Automated Configuration of Algorithms for Solving Hard Computational Problems. Ph.D. thesis, University Of British Columbia, Department of Computer Science, Vancouver, Canada (Oct. 2009).
- Hutter, F., Babić, D., Hoos, H.H. and Hu, A.J. Boosting verification by automatic tuning of decision procedures. In *Proceedings for Conference on Formal Methods in Computer-Aided Design* (2007), 27–34.
- Hutter, F., Hamadi, Y., Hoos, H.H., and Leyton-Brown, K. Performance prediction and automated tuning of randomized and parametric algorithms. In *Proceedings for Principles and Practice of Constraint Programming* (2006), 213–228.
- Hutter, F., Hoos, H.H. and Leyton-Brown, K. Automated configuration of mixed integer programming solvers. In *Proceedings for Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems* (2010), 186–202.
- Hutter, F., Hoos, H.H. and Leyton-Brown, K. Trade-offs in the empirical evaluation of competing algorithm designs. *Annals of Mathematics and Artificial Intelligence* 60, (2010), 65–89.
- Hutter, F., Hoos, H.H. and Leyton-Brown, K. Sequential model-based optimization for general algorithm configuration. In *Proceedings for Learning and Intelligent Optimization Conference* (2011), 507–523.
- Hutter, F., Tompkins, D.A.D. and Hoos, H.H. Scaling and probabilistic smoothing: Efficient dynamic local search for SAT. In *Proceedings for Principles and Practice of Constraint Programming* (2002), 233–248.
- Hutter, F., Xu, L., Hoos, H.H. and Leyton-Brown, K. Algorithm runtime prediction: Methods and evaluation. *Artificial Intelligence J.* 206 (Jan. 2014), 77–111.
- IBM. CPLEX Optimizer, 2014. <http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/>.
- Johnson, D.S. Random TSP generators for the DIMACS TSP Challenge, 2011; <http://dimacs.rutgers.edu/Challenges/TSP/>.
- KhudaBukhsh, A., Xu, L., Hoos, H.H. and Leyton-Brown, K. SATenstein: Automatically building local search SAT solvers from components. In *Proceedings for International Joint Conference on Artificial Intelligence* (2009), 517–524.
- Knuth, D. Estimating the efficiency of backtrack programs. *Mathematics of Computation* 29, 129 (1975), 121–136.
- Leyton-Brown, K., Nudelman, E., Andrew, G., McFadden, J. and Shoham, Y. Boosting as a metaphor for algorithm design. In *Proceedings for Principles and Practice of Constraint Programming* (2003), 899–903.
- Leyton-Brown, K., Nudelman, E., Andrew, G., McFadden, J. and Shoham, Y. A portfolio approach to algorithm selection. In *Proceedings for International Joint Conference on Artificial Intelligence* (2003), 1542–1543.
- Leyton-Brown, K., Nudelman, E. and Shoham, Y. Learning the empirical hardness of optimization problems: The case of combinatorial auctions. In *Proceedings for Principles and Practice of Constraint Programming* (2002), 556–572.
- Leyton-Brown, K., Nudelman, E. and Shoham, Y. Empirical hardness models: Methodology and a case study on combinatorial auctions. *Journal of the ACM* 56, 4 (2009), 1–52.
- Leyton-Brown, K., Pearson, M. and Shoham, Y. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings for ACM Conference on Electronic Commerce* (2000), 66–76.
- Mitchell, D., Selman, B. and Levesque, H. Hard and easy distributions of SAT problems. In *Proceedings for Conference on Artificial Intelligence* (1992), 459–465.
- Nudelman, E., Leyton-Brown, K., Andrew, G., Gomes, C., McFadden, J., Selman, B. and Shoham, Y. SATzilla 0.9. Solver description. SAT Competition, 2003.
- Nudelman, E., Leyton-Brown, K., Hoos, H.H., Devkar, A. and Shoham, Y. Understanding random SAT: Beyond the clauses-to-variables ratio. In *Proceedings for Principles and Practice of Constraint Programming* (2004), 438–452.
- Prasad, M.R., Biere, A. and Gupta, A. A survey of recent advances in SAT-based formal verification. *International Journal on Software Tools for Technology Transfer* 7, 2 (2005), 156–173.
- Rice, J.R. The algorithm selection problem. *Advances in Computers* 15 (1976), 65–118.
- Selman, B., Levesque, H.J. and Mitchell, D. A new method for solving hard satisfiability problems. In *Proceedings for Conference on Artificial Intelligence* (1992), 440–446.
- Vallati, M., Fawcett, C., Gerevini, A.E., Hoos, H.H. and Saetti, A. Generating fast domain-optimized planners by automatically configuring a generic parameterised planner. In *Proceedings for Automated Planning and Scheduling Workshop on Planning and Learning* (2011), 21–27.
- Xu, L., Hoos, H.H. and Leyton-Brown, K. Hierarchical hardness models for SAT. In *Proceedings for Principles and Practice of Constraint Programming* (2007), 696–711.
- Xu, L., Hoos, H.H. and Leyton-Brown, K. Predicting satisfiability at the phase transition. In *Proceedings for Conference on Artificial Intelligence* (2012), 584–590.
- Xu, L., Hutter, F., Hoos, H.H. and Leyton-Brown, K. SATzilla: Portfolio-based algorithm selection for SAT. *Journal of Artificial Intelligence Research* 32 (June 2008), 565–606.
- Xu, L., Hutter, F., Hoos, H.H. and Leyton-Brown, K. Evaluating component solver contributions to portfolio-based algorithm selectors. In *Proceedings for Theory and Applications of Satisfiability Testing* (2012), 228–241.
- Zarpas, E. Benchmarking SAT solvers for bounded model checking. In *Proceedings for Theory and Applications of Satisfiability Testing* (2005), 340–354.

**Kevin Leyton-Brown** (kevinlb@cs.ubc.ca) is an associate professor in the Department of Computer Science at the University of British Columbia, Canada.

**Holger H. Hoos** (hoos@cs.ubc.ca) is a professor in the Department of Computer Science at the University of British Columbia, Canada.

**Frank Hutter** is (fh@cs.uni-freiburg.de) Emmy Noether Research Group Lead, Department of Computer Science, at the University of Freiburg, Germany.

**Lin Xu** (xulin730@cs.ubc.ca) is a Ph.D. student in the Department of Computer Science at the University of British Columbia, Canada.

Copyright held by Owner/Author(s). Publication rights licensed to ACM. \$15.00.