

Run2Survive: A Decision-theoretic Approach to Algorithm Selection based on Survival Analysis

Alexander Tornede
Marcel Wever
Stefan Werner

Heinz Nixdorf Institute and Department of Computer Science, Paderborn University, Germany

ALEXANDER.TORNEDE@UPB.DE

MARCEL.WEVER@UPB.DE

STWERNER@MAIL.UPB.DE

Felix Mohr

Universidad de La Sabana, Chía, Cundinamarca, Colombia

FELIX.MOHR@UNISABANA.EDU.CO

Eyke Hüllermeier

Heinz Nixdorf Institute and Department of Computer Science, Paderborn University, Germany

EYKE@UPB.DE

Editors: Sinno Jialin Pan and Masashi Sugiyama

Abstract

Algorithm selection (AS) deals with the automatic selection of an algorithm from a fixed set of candidate algorithms most suitable for a specific instance of an algorithmic problem class, where “suitability” often refers to an algorithm’s runtime. Due to possibly extremely long runtimes of candidate algorithms, training data for algorithm selection models is usually generated under time constraints in the sense that not all algorithms are run to completion on all instances. Thus, training data usually comprises *censored* information, as the true runtime of algorithms timed out remains unknown. However, many standard AS approaches are not able to handle such information in a proper way. On the other side, survival analysis (SA) naturally supports censored data and offers appropriate ways to use such data for learning distributional models of algorithm runtime, as we demonstrate in this work. We leverage such models as a basis of a sophisticated decision-theoretic approach to algorithm selection, which we dub Run2Survive. Moreover, taking advantage of a framework of this kind, we advocate a risk-averse approach to algorithm selection, in which the avoidance of a timeout is given high priority. In an extensive experimental study with the standard benchmark ASlib, our approach is shown to be highly competitive and in many cases even superior to state-of-the-art AS approaches.

Keywords: algorithm selection, survival analysis, risk-aversion, surrogate loss

1. Introduction

Algorithm selection (AS) denotes the task of recommending algorithms for specific instances of an algorithmic problem class, such as the boolean satisfiability problem (SAT) or integer optimization. More specifically, the goal is to recommend the algorithm that is most suitable for the instance at hand, compared to other algorithms being also available. This idea is largely motivated by the observation of *performance complementarity* of algorithms, essentially suggesting that no algorithm dominates all others on all instances (Kerschke et al., 2019). In this regard, suitability is usually assessed in terms of a measure such as the quality of the solution returned, or the algorithm’s runtime in the case of constraint

satisfaction problems (CSP). Here, we focus on the latter, i.e., “being more suitable” is understood as having a shorter runtime.

In this context, a commonly encountered yet largely unsolved problem is the one of *censored training data*. Common AS approaches leverage observed runtime data together with machine learning methods such as regression (Xu et al., 2007; Amadini et al., 2014), (pairwise-)classification (Xu et al., 2011), or clustering (Kadioglu et al., 2010), to predict the runtime or the competitiveness of the algorithms on a new problem instance. However, the underlying data, consisting of runtimes measured for a specific algorithm on a certain problem instance, are often incomplete and partial: Not all algorithms have necessarily been run on all instances so far, and even if so, the observed runtimes might be *censored* due to a timeout. In the setting of AS, such timeouts are commonly set to save computational resources. The aforementioned approaches to runtime prediction have no way of directly dealing with censored data points, and instead apply ad-hoc solutions such as discarding them or imputing missing runtimes by default values. Most of these solutions are lacking a theoretical foundation, and, as will be detailed later on, some of them can lead to overoptimistic models.

On the other side, a well-established class of statistical methods summarized under the notion of *survival analysis* (Kleinbaum and Klein, 2010) is tailored to exactly such kind of censored data. Survival analysis deals with time-to-event data and seeks to model the underlying time-dependent phenomena. It originates from applications in medicine, where one is interested in the survival time of patients under various conditions. Luckily, it is not uncommon that a medical study ends before all patients die. Hence, the variable of interest, i.e., the survival time, can be censored in the sense of not being precisely observed. Methods in the field of survival analysis are able to handle such data in a proper way.

In this paper, we leverage survival analysis methods for the purpose of estimating the runtime of algorithms in the context of algorithm selection. More specifically, we make use of runtime data to estimate survival distributions that are specific to combinations of algorithms and problem instances. Since these are *generative* models providing rich information about runtimes in the form of fully specified probability distributions, they can serve as a basis of a sophisticated *decision-theoretic* approach to algorithm selection. Taking advantage of a framework of this kind, we advocate a *risk-averse* approach to algorithm selection, in which the avoidance of a timeout is given high priority.

2. The (Per-Instance) Algorithm Selection Problem

In Rice’s (per-instance) algorithm selection problem (Rice, 1976) we are faced with a problem instance space \mathcal{I} and a set of candidate algorithms \mathcal{A} , and we are interested in finding a mapping $s : \mathcal{I} \rightarrow \mathcal{A}$, called algorithm selector. For a given problem instance $i \in \mathcal{I}$, the algorithm selector is supposed to choose the algorithm $a^* \in \mathcal{A}$ optimizing a performance measure $m : \mathcal{I} \times \mathcal{A} \rightarrow \mathbb{R}$ which is costly to evaluate. Accordingly, the optimal selector, referred to as *oracle* or *virtual best solver* (VBS), is then defined as

$$s^*(i) = \arg \min_{a \in \mathcal{A}} \mathbb{E} [m(i, a)] \quad (1)$$

for all problem instances $i \in \mathcal{I}$, where the expectation accounts for the fact that $m(i, a)$ is a random variable due to possible randomness in the application of an algorithm.

2.1. Surrogate Models

Since m is costly to evaluate and the selection process itself typically underlies a time constraint, one cannot just evaluate m on all algorithms in \mathcal{A} . Therefore, at the core of the majority of existing AS approaches, machine learning models are employed for learning a surrogate model $\widehat{m} : \mathcal{I} \times \mathcal{A} \rightarrow \mathbb{R}$, which is fast to evaluate. Due to this property, an algorithm selector $s : \mathcal{I} \rightarrow \mathcal{A}$ can be specified as

$$s(i) := \arg \min_{a \in \mathcal{A}} \widehat{m}(i, a). \quad (2)$$

To infer a surrogate model, we are usually provided with a set of instances $\mathcal{I}_D \subset \mathcal{I}$ for which the performance in terms of m is known for *some but rarely all* algorithms $a \in \mathcal{A}$. Furthermore, as a prerequisite for generalizing runtime information across instances, we assume that these can be represented as d -dimensional feature vectors using a feature function $f : \mathcal{I} \rightarrow \mathbb{R}^d$. In the remainder of this paper, we write i for an instance and $\mathbf{i} = f(i)$ for its feature representation.

2.2. Censored Data

In combinatorial optimization, some algorithms may take extremely long to solve some instances (Gomes et al., 1997). Due to this, the generation of training data for surrogate models in algorithm selection is usually time-constrained in the sense that a time limit C , called *cutoff*, is set when running an algorithm. If the algorithm does not solve the instance prior to this timeout, the execution is aborted to save computational costs. In such a case, the runtime data is *right-censored* (Kleinbaum and Klein, 2010) in the sense that $m(i, a) > C$, i.e., the true runtime is known to exceed C . In cases where the data is not explicitly generated for training, but comes from a real-world scenario, it usually features only one true algorithm runtime for each instance, as in the absence of an algorithm selector, it is common practice to run many algorithms in parallel and stop all others as soon as the first one solves the instance. Hence, the data matrix $M \in \mathbb{R}^{|\mathcal{I}_D| \times |\mathcal{A}|}$ of known performance values spanned by the set of training instances \mathcal{I}_D and algorithms \mathcal{A} is only partially filled with known runtimes. As an example, consider the algorithm selection benchmark ASlib (Bischl et al., 2016), where in some cases more than 70% of the data points are censored (cf. Table 1).

A naïve approach to deal with this problem is to either impute missing runtimes with a default value or ignore them altogether when training a surrogate \widehat{m} . Common choices for a default are the cutoff time C or ten times the cutoff time, motivated by the *penalized average runtime* (PAR10) score, which is a common evaluation measure in the domain of AS. The PAR10 score corresponds to the runtime if the algorithm did not time out or ten times the cutoff if it timed out. However, all these strategies exhibit considerable drawbacks:

- Any form of imputation is a deliberate distortion of the training data and thus should be done with care. In scenarios with many censored data points, e.g., the one with over 70%, imputation may lead to strongly biased surrogate models. More specifically, an imputation of missing values with the cutoff time, which is lower than the actual runtime, will lead to a systematic underestimation of true runtimes.

- Dropping censored samples altogether is a waste of valuable information. Although the censored samples do not inform about precise runtimes, they still carry information, namely that $m(i, a) > C$. Furthermore, by dropping the long and keeping the short runtimes, there is again a danger of inducing overoptimistic models.

Although we are not the first to remark these problems of censored data (cf. Section 6), very few work has been done on solving them in the context of algorithm selection. A method for imputing censored data points introduced by Schmee and Hahn (1979) was studied in the context of algorithm configuration (AC) (Hutter et al., 2013; Eggenberger et al., 2018) and algorithm selection (Xu et al., 2007). A generalization of this method proposed by Eggenberger et al. (2018) starts by fitting a model on the uncensored data points, and then uses it to predict the mean μ and the variance σ^2 of the distribution for each censored data point. Based on these statistics, a truncated normal distribution $\mathcal{N}(\mu, \sigma^2)_{\geq C}$, where C is a known lower bound on the true runtime, is computed. Lastly, each censored data point is imputed with the mean of its associated truncated normal distribution and the model is refit based on both the censored and uncensored data points. This process is repeated until a stopping criterion is reached.

Despite the fact that this approach is a common way to deal with censored data in the context of AS and AC, the assumption of a normal distribution for runtimes is debatable and in contradiction to the more common heavy-tail assumption (Gomes et al., 1997). Moreover, the method was originally introduced for linear models and shown to work well for problems with only a single feature, but there is no strong justification for why the method should work well for higher-dimensional problems and more powerful models, too. In practice, one can observe that the method often fails to improve over training with the censored data directly (see Fig. 2).

Fortunately, survival analysis offers models which are able to *directly* work with censored data points in their learning procedure, without any need for imputation.

3. Survival Analysis and Random Survival Forests

In this section, we recall some basic concepts of survival analysis as developed in statistics. Moreover, we present a modern, non-parametric approach to fitting survival distributions, based on tree induction and ensembling techniques as commonly used in machine learning.

3.1. Basic Concepts of Survival Analysis

In survival analysis (SA), we typically proceed from historical data of the form

$$\mathcal{D} = \{(\mathbf{x}_n, y_n, \delta_n)\}_{n=1}^N, \quad (3)$$

where $\mathbf{x}_n \in \mathcal{X} \subseteq \mathbb{R}^d$ is a d -dimensional feature representation of a context, $y_n \in \mathbb{R}_+$ is the observed time until the event of interest occurred for the given context, and $\delta_n \in \{0, 1\}$ indicates whether the sample is (right-)censored ($\delta_n = 1$) or uncensored ($\delta_n = 0$), i.e., y_n is the true time until the event occurred or a clipped version thereof. More precisely,

$$y_n = \begin{cases} T_n & \text{if } \delta_n = 0 \\ C_n & \text{if } \delta_n = 1 \end{cases}, \quad (4)$$

where T_n is the uncensored survival time and C_n the censored one of instance n , which means that y_n is only a latent representation of T_n . Given such historical recordings, one is interested in inferring a model

$$z : \mathcal{X} \rightarrow \mathbb{R}_+, \tag{5}$$

which, given a context \mathbf{x} , correctly predicts the unknown survival time $T \in \mathbb{R}_+$.

In our setting of algorithm selection, the event of interest is the termination of an algorithm, and a context is given by a problem instance. Moreover, since we consider a fixed cutoff C for a timeout, $C_n = C$ for all $1 \leq n \leq N$.

Obviously, the problem of inducing (5) as stated above is very similar to the problem of standard regression. However, since the target variable is censored, the training data contains partial information and hence uncertainty. Accordingly, to capture this uncertainty, the majority of survival analysis approaches employs probabilistic models. To this end, the time until an event of interest T occurs is considered as a random variable, which is modeled via a probability distribution using the so-called *survival function* (SF):

$$S(t, \mathbf{x}) = \mathbb{P}[T \geq t \mid \mathbf{x}], \tag{6}$$

i.e., the probability that the event of interest occurs at time t or later, given context \mathbf{x} . Note that the survival function can be equivalently expressed in terms of the cumulative (death) distribution function (CDF) $F(t, \mathbf{x}) = \mathbb{P}[T \leq t \mid \mathbf{x}]$ as $S(t, \mathbf{x}) = 1 - F(t, \mathbf{x})$. While the survival function is defined according to the non-occurrence of the event until a certain time, the *hazard function*

$$h(t, \mathbf{x}) = \lim_{\Delta_t \rightarrow 0} \frac{\mathbb{P}[t \leq T < t + \Delta_t \mid T \geq t, \mathbf{x}]}{\Delta_t} \tag{7}$$

can be interpreted as expressing a degree of propensity of the event to occur at time t , under the condition that it did not occur before. The above functions are closely related to each other, and knowing one of them suffices to derive the others:

$$S(t, \mathbf{x}) = \exp\left(-\underbrace{\int_0^t h(u, \mathbf{x}) du}_{H(t, \mathbf{x})}\right), \quad h(t, \mathbf{x}) = -\left(\frac{dS(t, \mathbf{x})/dt}{S(t, \mathbf{x})}\right) \tag{8}$$

where $H(t, \mathbf{x})$ is the *cumulative hazard function*. To further understand this connection, consider the example of $S(t, \mathbf{x}) = \exp(-\lambda \cdot t)$ and $h(t, \mathbf{x}) \equiv \lambda$, where the hazard rate is modeled as a constant λ independent of the context. In practice, this assumption will of course be overly simplistic, and the hazard function will depend on the context features \mathbf{x} (also called *covariates* in SA). Moreover, to model the dependence on t , several parametric families of functions have been proposed, such as log-normal and Weibull functions.

A survival model can be used in various ways to obtain a real-valued prediction of the survival time, as requested by (5). A natural predictor is the expected survival time, i.e., the expectation of the random variable T :

$$z(\mathbf{x}) := \mathbb{E}[T] = \int_0^\infty S(t, \mathbf{x}) dt = \int_0^\infty 1 - F(t, \mathbf{x}) dt \tag{9}$$

3.2. Random Survival Forests

Due to their excellent predictive power, random forests are often chosen to tackle standard regression problems and also serve as a strong baseline in algorithm selection for modeling a surrogate model $\widehat{m}(i, a)$. For this reason, we chose to model the runtime distribution of an algorithm a via the survival function S_a in the form of a *random survival forest* (Ishwaran et al., 2008), an adaptation of standard random forests for survival analysis.

While being similar to standard random forests, random survival forests differ in the way they (a) build the individual survival trees and (b) generate predictions from these individual trees. Similar to the CART algorithm (Breiman et al., 1984), the individual survival trees are binary trees that are built via a recursive splitting approach. Splits are chosen to maximize survival difference between child nodes, i.e., by maximizing the difference in observed times y associated with these nodes. This process is continued until at least one of possibly multiple stopping criteria is reached, e.g., a node must contain at least d_0 uncensored samples. After such a survival tree has been built, it can be used to estimate the cumulative hazard function $H(t, \mathbf{x})$ for a query instance \mathbf{x} in the following way: Starting at the root node, one recursively determines which subtree the query instance \mathbf{x} belongs to depending on the split criterion of the considered node until a leaf node $e(\mathbf{x})$ is reached. Let $t_1 < \dots < t_{N(e(\mathbf{x}))}$ be the distinct times of events associated with node $e(\mathbf{x})$, $d_{t,e(\mathbf{x})}$ be the number of samples in node $e(\mathbf{x})$ where the event occurred at time t , i.e., with $\delta = 0$ and $y = t$, and $Y_{t,e(\mathbf{x})}$ the number of samples where the event has not occurred until time step t , i.e., with $\delta = 1$ or $y > t$. Then, the cumulative hazard function $H(t, \mathbf{x})$ is estimated using the Nelson-Aalen estimator as follows:

$$H(t, \mathbf{x}) = \sum_{t_i < \min(t_{N(e(\mathbf{x}))}, t)} \frac{d_{t_i, e(\mathbf{x})}}{Y_{t_i, e(\mathbf{x})}} \quad (10)$$

This can be seen as an empirical estimation of the hazard function for each event time t_i , and an accumulation of these estimates over time. Note that the cumulative hazard function fully defines the survival distribution for a single survival tree. To obtain the distribution based on the entire forest, one simply takes the mean over the cumulative hazard function estimates of the single trees.

4. Survival Analysis for Algorithm Selection

To tackle the problem of algorithm selection using SA, we learn one survival distribution, i.e., algorithm runtime distribution, for each algorithm a separately, using the above mentioned random survival forests. To this end, we leverage the training data available by constructing algorithm specific training datasets of the form (3) as

$$\mathcal{D}_a = \{ (i, m(i, a), \llbracket m(i, a) = C \rrbracket) \mid i \in \mathcal{I}_D \}, \quad (11)$$

where we assume that an occurrence of an algorithm runtime $m(i, a)$ equal to the cutoff time C , i.e., $m(i, a) = C$, indicates a timeout of algorithm a on instance i and hence a censored sample. Based on this, we estimate the associated survival distribution S_a for each algorithm a , which is fully defined by the associated cumulative hazard function H_a .

In practice, the random survival forests estimate survival *step* functions \widehat{S}_a for each algorithm a . Thus, we can approximate the integral of the survival function as follows:

$$\int_0^\infty S_a(t, \mathbf{i}) dt \approx \sum_{t_k, t_{k+1} \in \{0\} \cup \mathcal{T}_a \cup \{C\}} t_k \cdot \left[\widehat{S}_a(t_k, \mathbf{i}) - \widehat{S}_a(t_{k+1}, \mathbf{i}) \right], \quad (12)$$

where \mathcal{T}_a denotes the set of event times, i.e., points of termination, observed for algorithm a . Likewise, other integrals related to the survival function, such as the expected runtime, can be approximated.

4.1. Decision-theoretic Algorithm Selection

A survival distribution S_a can be seen as a generative model, which, in contrast to standard regression models, not only provides a point-estimate of an algorithm’s runtime, but characterizes the runtime of algorithm a on problem instance i in terms of a complete probability distribution of a random variable $T_{a,i}$. This is a rich source of information, which can be used to realize algorithm selection in a more sophisticated manner by means of a decision-theoretic approach, adopting the principle of *expected utility maximization* (Schoemaker, 1982), or, equivalently, *expected loss minimization*. More specifically, this principle suggests a very natural definition of the algorithm selector $s(\cdot)$ in (2):

$$s(i) := \arg \min_{a \in \mathcal{A}} \mathbb{E}[\mathcal{L}(T_{a,i})], \quad (13)$$

where $\mathcal{L} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a *loss function* which maps runtimes to real-valued loss degrees.

At first sight, a natural decision criterion is the expected survival time (9), i.e., the expected algorithm runtime, which was also suggested by Gagliolo and Schmidhuber (2006). This is obtained as a special case of (13) with $\mathcal{L}(t) = t$:

$$\mathbb{E}[\mathcal{L}(T_{a,i})] = \mathbb{E}[T_{a,i}] = \int_0^\infty S_a(t, \mathbf{i}) dt. \quad (14)$$

However, the expected algorithm runtime may appear sub-optimal in cases where the performance measure m of interest (cf. Sec. 2) substantially punishes algorithms running into a timeout. For example, the PAR10 score assigns the runtime of an algorithm as its score if it adheres to the timeout, but 10 times the cutoff if it times out. This is not accounted for by the expected runtime (14), which considers all regions of the survival function as *equally important* and tends to underestimate the *risk* of a timeout. In fact, an algorithm can have a shorter runtime than another *in expectation*, while having a higher probability of running into a timeout. To see this, consider the comparison between the first and third algorithm in Fig. 1 on some example instance; time is normalized to make the cutoff $C = 1$. The left plot shows the survival functions of five algorithms computed based on (8) and (10). For now, in the right plot, we focus only on the light blue curve, which, for each point in time t , depicts how much longer algorithm 1 is *expected* to run compared to algorithm 3, given that both algorithm run for at most t , i.e., $\mathbb{E}[T_1|T_1 \leq t] - \mathbb{E}[T_3|T_3 \leq t]$. Looking at the value of this curve at $t = 1$, it is clear that algorithm 3 has a better expected runtime (14) than algorithm 1. However, the left plot also shows that algorithm 3 has a substantially higher probability to time out, as its probability to terminate at $t = 1$, i.e., at the cutoff,

is larger than for algorithm 1. While a *risk-averse* view would focus on the survival probability at the timeout and then prefer algorithm 1 over algorithm 3, the expected algorithm runtime prefers 3 over 1, because it weights all parts of the distribution equally — clearly a sub-optimal choice when timeouts are strongly punished.

As an alternative, the PAR10 score itself may serve as a loss function:

$$\mathcal{L}(t) = \text{PAR10}(t) = \begin{cases} t & \text{if } t \leq C \\ 10 \cdot C & \text{else} \end{cases}. \quad (15)$$

The principle (13) of expected loss minimization then comes down to minimizing the *expected PAR10 score*. If this score plays the role of the *target loss*, on the basis of which the choice of an algorithm will eventually be assessed, this would indeed be a natural idea. It is all the more surprising that, to the best of our knowledge, this approach has not been considered in the literature so far. On the other side, it is true that the PAR10 loss is rather extreme in the sense of strongly focusing on the probability mass close to the cutoff. While we argued that uniformly averaging over the entire range of runtimes (i.e., computing the expected runtime) is sub-optimal, this behavior could be overly conservative and hence not be optimal either. To understand this, consider again Fig. 1. The desire to avoid timeouts at any cost would require to choose algorithm 1. At the same time, however, a preference for algorithm 3 appears by no means absurd, since there is undeniably a considerable probability that this algorithm is substantially faster than algorithm 1 — all the more keeping in mind that the survival functions are only estimates of the underlying ground truth distributions, and hence $\mathbb{E}[\mathcal{L}(T_{a,i})]$ only an estimate of the true expected loss.

4.2. Risk-Averse Algorithm Selection

In light of the above discussion, our general goal should be to trade off the different regions of the runtime distributions, keeping in mind the advantages of potentially very short runtimes on the one hand and the risk of timeouts on the other hand. To this end, we consider two parameterized classes of functions for instantiating \mathcal{L} . These functions can be seen as *surrogate losses*, i.e., surrogates of the actual target loss (PAR10). Both functions are convex, thereby reflecting a *risk-averse* attitude in decision making (Fishburn, 1969). In the context of algorithm selection, where decisions are algorithms and the avoidance of timeouts has highest priority, risk aversion appears to be a very natural and meaningful property.

The first class of functions consists of polynomials

$$\mathcal{L}_\alpha(t) = t^\alpha, \quad (16)$$

where $\alpha \in \mathbb{R}_+$ controls the degree of risk aversion. Hence, larger values of $\alpha \geq 1$ result in stronger risk aversion and are therefore better suited for hard algorithmic problems with a large risk of timeouts, whereas smaller values $0 < \alpha < 1$ should be chosen for simpler problems where the majority of algorithms terminates before the cutoff time. One can observe that the higher α , the less important is the behavior of the algorithms for low runtimes (the head of the distribution) and the more important for long runtimes (the tail of the distribution), leading to a clear preference of the “safer” alternatives for $\alpha > 2$. In our experimental evaluation, we demonstrate that α can be tuned effectively using hyperparameter optimization techniques.

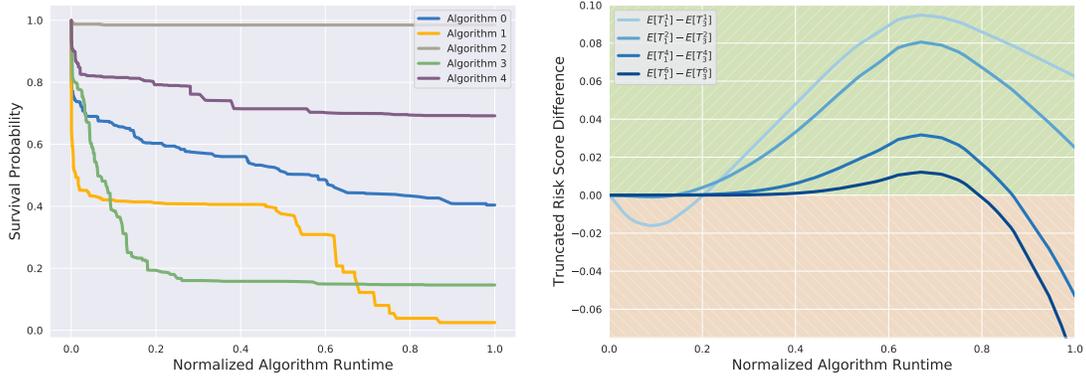


Figure 1: Left: Survival functions of various algorithms on a problem instance. Right: Truncated risk score difference between the runtime of algorithm 1 (yellow on the left) and algorithm 3 (green), i.e. $\mathbb{E}[T_1^\alpha | T_1 \leq t] - \mathbb{E}[T_3^\alpha | T_3 \leq t]$. Values below the zero line for the non-truncated score indicate that $\mathbb{E}[T_1^\alpha] < \mathbb{E}[T_3^\alpha]$ and hence algorithm 1 is selected over algorithm 3, which is only the case with higher risk aversion. Furthermore, indeed larger values of α emphasize the tail of the distributions as the difference is close to the zero for small values of t .

As a second class of functions, we consider surrogate losses of the form

$$\mathcal{L}_{\alpha,\beta}(t) = \min \{-\alpha \log(1 - t), \beta\}, \quad (17)$$

where $\alpha \in [0, 1]$ again defines the degree of risk aversion, with smaller values encouraging more risk-averse selections, i.e., potentially safer algorithm choices, and $\beta \in \mathbb{R}_+$ constitutes an upper bound on \mathcal{L} to limit extreme behavior. Here, we assume that $t \in [0, 1]$, which can be achieved via rescaling all runtimes such that $C = 1$. Again, both parameters α and β can be learned as we demonstrate in the experiments.

In practice, choosing the right surrogate loss and optimally adapting its parameters to the problem at hand is a difficult task. Therefore, we suggest to determine the most suitable surrogate loss in a data-driven way using hyperparameter optimization techniques. In the experimental evaluation, we present an approach based on Bayesian optimization (Frazier, 2018), automatically selecting either the polynomial or the log-based surrogate loss as well as adequate parameters α (and β).

5. Experimental Evaluation

In this section, we provide an extensive evaluation of the proposed methodology based on survival analysis, applying it to the algorithm selection library ASlib v4.0 benchmark suite (Bischl et al., 2016) and comparing it to state-of-the-art approaches. The benchmark suite comprises a broad variety of scenarios from the literature beyond commonly addressed AS domains. A comprehensive overview of the considered scenarios is provided in Table 1.

Table 1: Overview of examined ASlib scenarios including their number of instances (#I), unsolved instances (#U), algorithms (#A) and provided features (#F). We also show the cutoffs (C) and the percentage of censored algorithm runs (%C).

	ASF-POTASSCO	BNSL-2016	CPMP-2015	CSF-2010	CSP-MZN-2013	CSP-Minimize-Time-2016	GRAPHS-2015	MAXSAT-PMS-2016	MAXSAT-WPMS-2016	MAXSAT12-PMS	MAXSAT15-PMS-INDU	MIP-2016	PROTEUS-2014	QBF-2011	QBF-2014	QBF-2016	SAT03-16-INDU	SAT11-HAND	SAT11-INDU	SAT11-RAND	SAT12-ALL	SAT12-HAND	SAT12-INDU	SAT12-RAND	SAT15-INDU	TSP-LION2015
#I	1294	1179	527	2024	4642	100	5725	601	630	876	601	218	4021	1368	1254	825	2000	296	300	600	1614	767	1167	1362	300	3106
#U	82	0	0	253	944	17	117	45	89	129	44	0	456	314	241	55	269	77	47	108	20	229	209	322	17	0
#A	11	8	4	2	11	20	7	19	18	6	29	5	22	5	14	24	10	15	18	9	31	31	31	31	28	4
#F	138	86	22	86	155	95	35	37	37	37	37	143	198	46	46	46	483	115	115	115	115	115	115	115	54	122
C	600	7200	3600	5000	1800	1200	1e+08	1800	1800	2100	1800	7200	3600	3600	900	1800	5000	5000	5000	5000	1200	1200	1200	1200	3600	3600
%C	20.1	28.1	27.6	19.6	70.2	50	7.0	39.4	57.9	41.2	48.9	20	60.2	54.7	55.6	36.2	24.7	60.7	33.3	47.5	53.9	66.7	50.4	73.5	23.5	9.6

5.1. Experimental Setup

We evaluate each approach conducting a 10-fold cross-validation for every scenario. Unsolvable instances, for which no valid solution could be determined by any algorithm before violating the fixed timeout, are subsequently removed from the test fold.

More sophisticated industrial algorithm selection systems commonly integrate complementary techniques such as pre-solvers or feature selectors. For our experimental study, we consider pure algorithm selection models without any pre-solvers or feature selectors to focus on the models themselves and their suitability to deal with censored data.

The performance of the approaches is assessed via the PAR10 metric defined as the penalized average runtime of the selected solver required to solve problem instances. Here, timeout violations are explicitly penalized by a factor of 10. Whether the proposed techniques constitute improvements over the Single Best Solver (SBS) baseline, and consequently close the gap to the Virtual Best Solver, can be assessed in terms of the *normalized PAR10* (nPAR10) metric defined as $nPAR10 = (PAR10_{Model} - PAR10_{VBS}) / (PAR10_{SBS} - PAR10_{VBS})$. Here, the SBS refers to the algorithm being best on average wrt. PAR10, whereas the (hypothetical) VBS denotes the optimal algorithm selector. Normalized scores greater than 1 denote that the approach is on average less efficient than the SBS, whereas scores below 1 indicate an improvement over the SBS baseline.

The time for computing the corresponding instance feature representations required by the approaches is accounted for as well.

Our experimental study specifically distinguishes the examined baselines according to their performance under different treatment of censored data as discussed in Section 2.2. We initially examine which of the previously proposed methods to address censoring complies best with each baseline, and subsequently compare their respectively best results in terms of their PAR10 performance against our proposed methodology. Consequently, each baseline selector is evaluated under removal of censored data (*Ignored*), imputation of the cutoff (*Runtime*), and imputation of timeout penalty (*PAR10*). Where applicable, we also evaluate the method proposed by Schmee and Hahn (1979). For each baseline selector, the most effective method to cope with censoring is assessed in terms of median PAR10 performance aggregated over all studied ASlib scenarios.

In the overall comparison, we include three approaches leveraging random survival forests for estimating the survival functions. First, we consider an approach selecting algorithms according to their expected runtime as defined in (14) (referred to as Run2SurviveExp). Second, a selection according to the expected PAR10 score of an algorithm as defined in (15) (referred to as Run2SurvivePAR10) is evaluated. Lastly, a method making use of the polynomial and log-based surrogate loss functions is used, where (a) the type of the surrogate loss and (b) respective parameters are tuned by means of Bayesian optimization on the training data (Run2SurvivePoly/Log).

All experiments were run on machines featuring Intel Xeon E5-2695v4@2.1GHz CPUs with 16 cores and 64GB RAM. Moreover, the code for all experiments and the generation of plots including detailed documentation can be found on Github¹.

5.2. Baselines

In the following, we detail all baselines we compare with.

SUNNY (Amadini et al., 2014) retrieves the k nearest neighbours of a queried problem instance in terms of Euclidean distance w.r.t. the feature representation. From this set, the most efficient algorithm in terms of the (possibly imputed) training runtimes is chosen.

ISAC (Kadioglu et al., 2010) leverages a g-means clustering algorithm instead of a kNN approach to partition the feature space. Originally designed for algorithm configuration, we transfer ISAC to algorithm selection borrowing SUNNY’s decision strategy. More specifically, we first assign the new instance to the closest centroid and subsequently select algorithms performing best on the training instances associated with the respective cluster.

SATzilla’11 In contrast, SATzilla’11 (Xu et al., 2011) leverages pairwise comparisons of algorithms. For each pair of algorithms, a Decision Forest is trained to decide which one performs better for a given problem instance. An explicit cost-sensitive loss function weights instances subject to the absolute difference in their imposed computational cost. Faced with a new instance, the algorithm that wins the highest number of pairwise comparisons is predicted.

PerAlgorithmRegressor The PerAlgorithmRegressor learns one Random Forest per candidate algorithm to predict each algorithm’s performance while facilitating algorithm selection according to their respective estimates.

MultiClassSelector The MultiClassSelector baseline follows a multi-class classification approach, where instances are first labeled with the most effective selections, and a Random Forest is subsequently employed as classifier to establish suitable selections.

5.3. Results

Fig. 2 summarizes the results for the examined baselines subject to different methods of handling censoring. We find that imputing censored algorithm runs with their timeout works best for the PerAlgorithmRegressor, SATzilla’11, and SUNNY, whereas the PAR10 imputation is more effective for ISAC. We further find an iterative adjustment of censored values according to the technique proposed by Schmees and Hahn (1979) to not benefit the PerAlgorithmRegressor in terms of median performance. Since an instance’s most efficient

1. https://github.com/alexandertornede/algorithm_survival_analysis

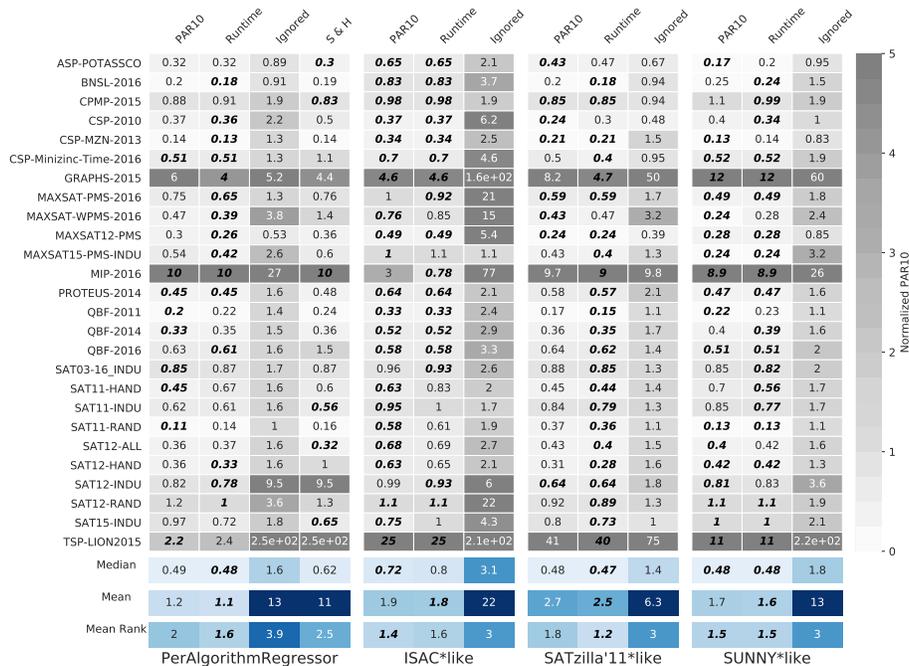


Figure 2: Normalized PAR10 results of baselines for different ways of dealing with censored data: labeling data points as proposed by Schmee and Hahn (1979) (S&H), with the PAR10 score (PAR10), the cutoff C (runtime), or the corresponding data points are ignored. The best results for each scenario are printed in bold.

solver is preserved regardless of whether censored algorithm runs are imputed with their exact timeout, PAR10 penalty or even entirely ignored, we note that the MultiClassSelector’s performance is insensitive to prior imputation methods and thus not included in Fig. 2. While treating censored values as their exact timeout or PAR10 score evidently yields similar results, obviously, discarding censored algorithm runs is the least appealing method, as depending on the ASlib scenario up to 73.5% of runs are ignored.

Fig. 3 compares the methods based on survival analysis proposed in this paper against the respectively most effective results w.r.t. the imputation procedure for each previously discussed baseline. Best results for each scenario are printed in bold, and an overline indicates beating all baselines. Evidently, neither ISAC nor the MultiClassSelector establish competitive algorithm selection on the ASlib benchmark, as the former approach only constitutes an effective choice for one sole scenario, whereas the latter is consistently outperformed. The PerAlgorithmRegressor, SUNNY and SATzilla’11, in contrast, broadly attain competitive results and consequently represent strong competitors to the random survival forests. However, Fig. 3 illustrates each of our Run2Survive variants to outperform their adversaries in terms of median, mean nPAR10, as well as mean rank performance aggregated across all scenarios. Compared to the baselines, Run2SurviveExp facilitates more effective per-instance algorithm selection on 8 scenarios, whereas Run2SurvivePAR10 yields superior results on 13 scenarios. While Run2SurvivePoly/Log beats all baselines in only 12 scenar-

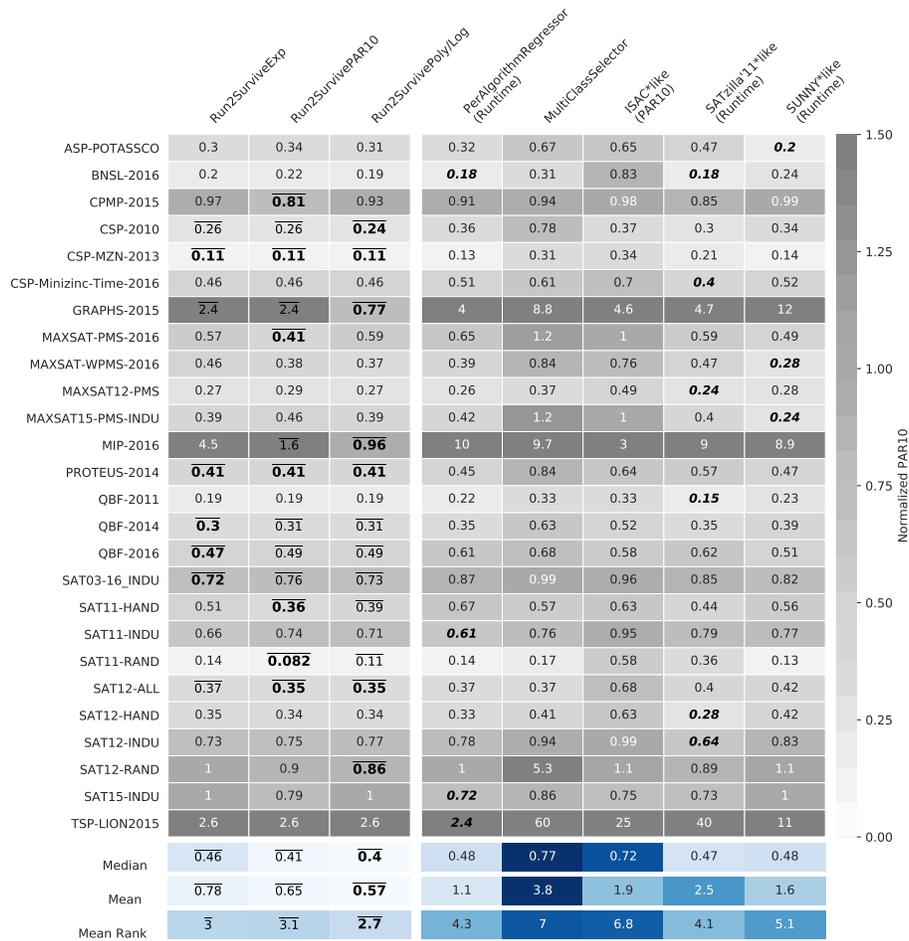


Figure 3: Normalized PAR10 results where for each baseline the way of dealing with censored data is selected according to the minimum median across all examined scenarios. The best results for each scenario are printed in bold whereas an overline indicates beating all baselines.

ios, it achieves an nPAR10 score below 1 on all scenarios except one, and hence consistently beats the SBS, which is not the case for any baseline. Furthermore, Run2SurvivePoly/Log beats the SBS in two scenarios where no other approach is able to.

Our findings further illustrate the usefulness of risk-averse decision making in algorithm selection, as Run2SurviveExp can be improved upon in terms of median and mean nPAR10 performance, and also in mean rank (Run2SurvivePoly/Log). In detail, Run2SurvivePAR10 and Run2SurvivePoly/Log identify superior selection criteria in 10 and 11 ASlib scenarios. However, they also degrade compared to Run2SurviveExp in 9 respectively 7 cases.

6. Related Work

In this section, we give a brief overview of closely related work. For a more comprehensive overview of algorithm selection and runtime prediction in general, we refer to (Kerschke et al., 2019) and (Hutter et al., 2014), respectively.

To the best of our knowledge, the first and only work making direct use of survival analysis in the context of algorithm selection is (Gagliolo and Schmidhuber, 2006; Gagliolo and Legrand, 2010). Gagliolo and Schmidhuber (2006) present an algorithm called GambleTA computing adaptive algorithm schedules for SAT and the Auction Winner Determination problem instances without any prior learning phase. GambleTA is a bandit-based algorithm deciding online which algorithms to run for how long such that it can use the obtained data to learn the runtime distributions of the considered algorithms. These distributions are modeled using cumulative hazard functions (CHF), which in turn are used to compute static schedules based on different techniques, such as selecting at timestep t the algorithm with e.g. minimal expected runtime. Using these static schedules, the authors also show how to construct dynamic schedules.

Our work distinguishes itself from theirs mainly by the following points: First, Gagliolo and Schmidhuber (2006) consider the setting of dynamic algorithm scheduling with a need for online learning, whereas we consider the standard setting of algorithm selection, whence other recommendations are required. Second, while the authors only consider two algorithmic problems and algorithm sets of size two, we consider the complete ASlib with over 25 scenarios from different problem domains featuring between 2 and 31 algorithms, making our experimental study much more extensive. Third, the authors consider nearest-neighbor motivated survival analysis models for modeling the CHF of the algorithms, while we make use of random survival forests (although in principle any technique could be used). Finally, and perhaps most importantly, we mainly focus on designing a theoretically sound algorithm selection strategy making use of the concept of risk aversion, leveraging surrogate loss functions to better tailor selections towards the structure of loss functions similar to the PAR10 score.

Xu et al. (2007) also remark the problems and implications of censored data in the context of algorithm runtime prediction. Instead of directly applying a method tailored towards such data, they apply the method by Schmee and Hahn (1979) discussed in Section 2.2. Furthermore, similar to us, they perform a case study on how treating censored data during the training of runtime models impacts these models' performance. They find that (1) both ignoring and imputing censored data with the cutoff time does indeed yield overly optimistic models (as we also note in Section 2.2), and (2) that the method proposed by Schmee and Hahn (1979) yields best performance in terms of the RMSE.

In the context of algorithm configuration, Hutter et al. (2013, 2014) also note problems arising from censored data when not treating these datapoints appropriately and describe a generalization of the method of Schmee and Hahn (1979) to random forests making these better suited for censored data. Similarly, Biedenkapp et al. (2017) and Eggenberger et al. (2018) shortly mention the problem of censored data arising from terminating all running algorithm configurations when one out of many parallel runs finishes, and suggest to impute such censored data by previously discussed means.

7. Conclusion

In this paper, we proposed the use of survival analysis techniques combined with a decision-theoretic approach for the problem of algorithm selection. Taking advantage of the rich information provided by generative models of algorithm runtimes, together with the use of a risk-averse decision strategy to select the most promising algorithm for an unseen instance, we achieved a robust overall performance across different problem domains, such as SAT, CSP, and CPMP. Applying our method to a suite of 26 benchmark scenarios for algorithm selection from the standard benchmark library ASlib, we find it to be highly competitive and in many cases even superior to state-of-the-art methods. Moreover, considering several statistics across the considered benchmark datasets, our approach performs best in terms of average rank as well as median/mean normalized PAR10 score across all scenarios, while achieving best performances in up to 13 out of 26 scenarios.

Encouraged by the strong empirical results, we plan to further elaborate on our methodology in future work. For example, it would be tempting to transfer it to the extreme algorithm selection (XAS) scenario (Tornede et al., 2019, 2020), which deals with a significantly larger collection of algorithms to choose from—correspondingly, XAS requires a feature representation not only for problem instances but also for algorithms. Another interesting direction is the integration of Run2Survive with a combined ranking and regression approach (Hanselle et al., 2020).

Acknowledgments

This work was partially supported by the German Research Foundation (DFG) within the Collaborative Research Center “On-The-Fly Computing” (SFB 901/3 project no. 160364472). The authors also gratefully acknowledge support of this project through computing time provided by the Paderborn Center for Parallel Computing (PC²).

References

- R. Amadini, M. Gabbrielli, and J. Mauro. SUNNY: a lazy portfolio approach for constraint solving. *Theory Pract. Log. Program.*, 14(4-5), 2014.
- A. Biedenkapp, M. Lindauer, K. Eggenberger, F. Hutter, C. Fawcett, and H. H. Hoos. Efficient parameter importance analysis via ablation with surrogates. In *AAAI*, 2017.
- B. Bischl, P. Kerschke, L. Kotthoff, M. Lindauer, Y. Malitsky, A. Fréchette, H. H. Hoos, F. Hutter, K. Leyton-Brown, K. Tierney, and J. Vanschoren. ASlib: A benchmark library for algorithm selection. *Artif. Intell.*, 237, 2016.
- L. Breiman, J. Friedman, C. J Stone, and R. A Olshen. *Classification and regression trees*. CRC press, 1984.
- K. Eggenberger, M. Lindauer, H. H. Hoos, F. Hutter, and K. Leyton-Brown. Efficient benchmarking of algorithm configurators via model-based surrogates. *Mach. Learn.*, 107(1), 2018.
- P. C Fishburn. *Utility-Theory for Decision Making*. Wiley, 1969.

- P. I. Frazier. A tutorial on bayesian optimization. *CoRR*, abs/1807.02811, 2018.
- M. Gagliolo and C. Legrand. Algorithm survival analysis. In *Experimental Methods for the Analysis of Optimization Algorithms*. Springer, 2010.
- M. Gagliolo and J. Schmidhuber. Learning dynamic algorithm portfolios. *Ann. Math. Artif. Intell.*, 47(3-4), 2006.
- C. P. Gomes, B. Selman, and N. Crato. Heavy-tailed distributions in combinatorial search. In *CP*, 1997.
- J. Hanselle, A. Tornede, M. Wever, and E. Hüllermeier. Hybrid ranking and regression for algorithm selection. In *KI 2020: Advances in Artificial Intelligence*, 2020.
- F. Hutter, H. H. Hoos, and K. Leyton-Brown. Bayesian optimization with censored response data. *CoRR*, abs/1310.1947, 2013.
- F. Hutter, L. Xu, H. H. Hoos, and K. Leyton-Brown. Algorithm runtime prediction: Methods & evaluation. *Artif. Intell.*, 206, 2014.
- H. Ishwaran, U. B Kogalur, E. H Blackstone, M. S Lauer, et al. Random survival forests. *The annals of applied statistics*, 2(3), 2008.
- S. Kadioglu, Y. Malitsky, M. Sellmann, and K. Tierney. ISAC - instance-specific algorithm configuration. In *ECAI*, 2010.
- P. Kerschke, H. H. Hoos, F. Neumann, and H. Trautmann. Automated algorithm selection: Survey and perspectives. *Evol. Comput.*, 27(1), 2019.
- D. G Kleinbaum and M. Klein. *Survival analysis*, volume 3. Springer, 2010.
- J. R. Rice. The algorithm selection problem. *Advances in Computers*, 15, 1976.
- J. Schmee and G. J Hahn. A simple method for regression analysis with censored data. *Technometrics*, 21(4), 1979.
- P.J. Schoemaker. The expected utility model: Its variations, purposes, evidence and limitations. *Journal of Economic Literature*, 20:529–563, 1982.
- A. Tornede, M. Wever, and E. Hüllermeier. Algorithm selection as recommendation: From collaborative filtering to dyad ranking. In *CI Workshop, Dortmund*, 2019.
- A. Tornede, M. Wever, and E. Hüllermeier. Extreme algorithm selection with dyadic feature representation. In *Discovery Science*, 2020.
- L. Xu, F. Hutter, H. H Hoos, and K. Leyton-Brown. Satzilla-07: the design and analysis of an algorithm portfolio for sat. In *CP*. Springer, 2007.
- L. Xu, F. Hutter, H. Hoos, and K. Leyton-Brown. Hydra-mip: Automated algorithm configuration and selection for mixed integer programming. In *RCRA workshop @ IJCAI*, 2011.