



HAL
open science

Sampling Effect on Performance Prediction of Configurable Systems: A Case Study

Juliana Alves Pereira, Mathieu Acher, Hugo Martin, Jean-Marc Jézéquel

► **To cite this version:**

Juliana Alves Pereira, Mathieu Acher, Hugo Martin, Jean-Marc Jézéquel. Sampling Effect on Performance Prediction of Configurable Systems: A Case Study. 2019. hal-02356290v1

HAL Id: hal-02356290

<https://inria.hal.science/hal-02356290v1>

Preprint submitted on 29 Nov 2019 (v1), last revised 21 Apr 2020 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sampling Effect on Performance Prediction of Configurable Systems: A Case Study

JULIANA ALVES PEREIRA, Univ Rennes, Inria, CNRS, IRISA

MATHIEU ACHER, Univ Rennes, Inria, CNRS, IRISA

HUGO MARTIN, Univ Rennes, Inria, CNRS, IRISA

JEAN-MARC JÉZÉQUEL, Univ Rennes, Inria, CNRS, IRISA

Numerous software systems are highly configurable and provide a myriad of configuration options that users can tune to fit their functional and performance requirements (e.g., execution time). Measuring all configurations of a system is the most obvious way to understand the effect of options and their interactions, but is too costly or infeasible in practice. Numerous works thus propose to measure only a few configurations (a sample) to learn and predict the performance of any combination of options' values. A challenging issue is to sample a small and representative set of configurations that leads to a good accuracy of performance prediction models. A recent study devised a new algorithm, called distance-based sampling, that obtains state-of-the-art accurate performance predictions on different subject systems. In this paper, we replicate this study through an in-depth analysis of x264, a popular and configurable video encoder. We systematically measure 1,152 configurations of x264 with 17 input videos and two quantitative properties (encoding time and encoding size). Our goal is to understand whether there is a dominant sampling strategy over the very same subject system (x264), i.e., whatever the workload and targeted performance properties. The findings from this study show that random sampling leads to more accurate performance models. However, without considering random, there is no single "dominant" sampling, instead different strategies perform best on different inputs and non-functional properties, further challenging practitioners and researchers.

Additional Key Words and Phrases: Software Product Lines, Configurable Systems, Machine Learning, Performance Prediction

1 INTRODUCTION

Configurable software systems offer a multitude of configuration options that can be combined to tailor the systems' functional behavior and performance (e.g., execution time, memory consumption). Options often have a significant influence on performance properties that are hard to know and model *a priori*. There are numerous possible options values, logical constraints between options, and subtle interactions among options [14, 24, 45, 50, 51] that can have an effect while quantitative properties such as execution time are themselves challenging to comprehend.

Measuring all configurations of a configurable system is the most obvious path to e.g., find a well-suited configuration, but is too costly or infeasible in practice. Machine-learning techniques address this issue by measuring only a subset of configurations (known as sample) and then using these configurations' measurements to build a performance model capable of predicting the performance of other configurations (i.e., configurations not measured before). Several works thus follow a "sampling, measuring, learning" process [2, 14, 19–22, 26, 39, 41, 42, 45, 50, 51, 59, 62, 63]. A crucial step is

the way the sampling is realized, since it can drastically affect the performance model accuracy [24, 45]. Ideally, the sample should be small to reduce the cost of measurements and representative of the configuration space to reduce prediction errors. The sampling phase involves a number of difficult activities: (1) picking configurations that are valid and conform to constraints among options – one needs to resolve a satisfiability problem; (2) instrumenting the executions and observations of software for a variety of configurations – it might have a high computational cost especially when measuring performance aspects of software; (3) guaranteeing a coverage of the configuration space to obtain a representative sample set. An ideal coverage includes all influential configuration options by covering different kinds of interactions relevant to performance. Otherwise, the learning may hardly generalize to the whole configuration space.

With the promise to select a small and representative sample set of valid configurations, several sampling strategies have been devised in the last years [45]. For example, random sampling aims to cover the configuration space uniformly while coverage-oriented sampling selects the sample set according to a coverage criterion (e.g., *t*-wise sampling to cover all combinations of *t* selected options). Recently, Kaltenecker *et al.* [24] analyzed 10 popular real-world software systems and found that their novel proposed sampling strategy, called diversified distance-based sampling, dominates five other sampling strategies by decreasing the cost of labelling software configurations while minimizing the prediction errors.

In this paper, we conduct a replication of this preliminary study by exclusively considering the x264 case, a configurable video encoder. Though we only consider one particular configurable system, we make vary its workloads (with the use of 17 input videos) and measured two performance properties (encoding time and encoding size) over 1,152 configurations. Interestingly, Kaltenecker *et al.* [24] also analyzed the same 1,152 configurations of x264, but a fixed input video was used and only the time was considered. The goal of our experiments is to determine whether sampling strategies considered in [24] over different subject systems are as effective on the *same* configurable system, but with different factors possibly influencing the distribution of the configuration space. An hypothesis is that practitioners of a configurable system can rely on a one-size-fits-all sampling strategy that is cost-effective whatever the factors influencing the distribution of its configuration space. On the contrary, another hypothesis is that practitioners should change their sampling strategy each time an input (here: videos) or a performance property are targeted.

We investigate to what extent sampling strategies are sensitive to different workloads of the x264 configurable system and different

Authors' addresses: Juliana Alves Pereira, Univ Rennes, Inria, CNRS, IRISA, Rennes, France, juliana.alves-pereira@irisa.fr; Mathieu Acher, Univ Rennes, Inria, CNRS, IRISA, Rennes, France, mathieu.acher@irisa.fr; Hugo Martin, Univ Rennes, Inria, CNRS, IRISA, Rennes, France, hugo.martin@irisa.fr; Jean-Marc Jézéquel, Univ Rennes, Inria, CNRS, IRISA, Rennes, France, jean-marc.jezequel@irisa.fr.

performance properties: What are the most effective sampling strategies? Is random sampling a strong baseline? Is there a dominant sampling strategy that practitioners can always rely on? To this end, we systematically report the influence of sampling strategies on the accuracy of performance predictions and on the robustness of prediction accuracy.

Our contributions are as follows:

- We rank six sampling strategies based on a unified benchmark over seventeen different inputs and two non-functional properties. We find that uniform random sampling dominates all the others sampling strategies w.r.t. prediction accuracy for a majority of cases;
- We show that the ranking of non-random sampling strategies is quite unstable, being heavily sensitive to input video and targeted performance property. We gather preliminary insights about the effectiveness of some sampling strategies on some videos and performances;
- We compare our results to the previous study of x264 [24], that was based on the use of a single input video and measurements of encoding time. Through our experimental results, we find that there is not a single “dominant” approach as in Kaltenecker et al. [24];
- We have made the data and code from our replication study publicly available at <https://github.com/FAMILIAR-project/x264-inputsensitivity>.

With respect to the categorized research methods by Stol et al. [56], our paper mainly contributes to a knowledge-seeking study. Specifically, we perform a field experiment of x264, a highly-configurable system and mature project. We gain insights about performance’ properties using a large corpus of configurations and input videos.

Audience. Researchers and practitioners in configurable systems and performance engineering shall benefit from our sampling experiments and insights. We discuss impacts of our results for practitioners (How to choose a sampling strategy?) and for the research community (e.g., on the design and assessment of sampling strategies).

2 BACKGROUND AND RELATED WORK

In this section, we introduce basic concepts of configurable software systems and motivate the use of learning techniques in this field. Furthermore, we briefly describe six state-of-the-art sampling strategies used in our experiments.

2.1 Learning Software Configuration Spaces

x264 is a command-line tool to encode video streams into the H.264/MPEG-4 AVC format. Users can configure x264 through the selection of numerous options, some having an effect on the time needed to encode a video, on the quality or the size of the output video, etc. A configuration of x264 is an assignment of values to options. In our study and as in [24], we only consider Boolean options that can be selected or deselected. As in most configurable systems, not all combinations of options are valid due to constraints among options. For instance, `ref_1`, `ref_5`, and `ref_9` are mutually exclusive and at least one of these options should be selected.

Executing and measuring every valid configuration to know about its performance or identify the performance-optimal one is often unfeasible or costly. To overcome this problem, machine learning

techniques rely on a small and representative *sample* of configurations (see Figure 1). Each configuration of the sample is executed and labelled with a performance measurement (e.g., encoding time). The sample is then used for *training* a learning algorithm (i.e., a regressor) that builds a performance model capable of predicting the performance of unmeasured configurations. The performance model may lead to prediction errors. The overall goal is to obtain high *accuracy* roughly computed as the difference between actual performances and predicted performances (more details are given hereafter). How to efficiently sample, measure, and learn is subject to intensive research [45]. In case important (interactions among) options are not included in the training set, the learning phase can hardly generalize to the whole population of configurations. Hence, sampling is a crucial step of the overall learning process with an effect on the accuracy of the prediction model.

2.2 Sampling Strategies

Several sampling strategies have been proposed in the literature about software product lines and configurable systems [45, 61, 65].

Sampling for testing. Some works consider sampling for the specific case of testing configurations. There is not necessarily a learning phase and the goal of sampling is mostly to find and cover as many faults as possible. For instance, Medeiros et al. compared 10 sampling algorithms to detect different faults in configurable systems [34]. Arcuri et al. theoretically demonstrate that a uniform random sampling strategy may outperform coverage-based sampling [4] (see hereafter). Halin et al. demonstrated that uniform random sampling forms a strong baseline for faults and failure efficiency on the JHipster case [16]. Varshosaz et al. [65] conducted a survey of sampling for testing configurable systems. Though the purpose differs, some of these sampling strategies are also relevant and considered in the context of performance prediction.

Sampling for learning. Pereira et al. [45] review several sampling strategies specifically used for learning configuration spaces. We now present an overview of six sampling strategies also considered in [24] and used in our study. All strategies have the merit of being agnostic of the domain (no specific knowledge or prior analysis are needed) and are directly applicable to any configurable system.

Random. sampling aims to cover the configuration space *uniformly*. Throughout the paper, we refer to random as uniform random sampling. The challenge is to select one configuration amongst all the *valid* ones in such a way each configuration receives an equal probability to be included in the sample. An obvious solution is to enumerate all valid configurations and randomly pick a sample from the whole population. However, enumerative approaches quickly do not scale with a large number of configurations. Oh et al. [42] rely on binary decision diagrams to compactly represent a configuration space, which may not scale for very large systems [35]. Another line of research is to rely on satisfiability (SAT) solvers. For instance, UniGen [7, 8] uses a hashing-based functions to synthesize samples in a nearly uniform manner with strong theoretical guarantees. These theoretical properties come at a cost: the hashing-based approach requires adding large clauses to formulas. Plazar et al. [46] showed that state-of-the-art algorithms are either not able to produce any sample or unable to generate uniform samples for the SAT instances

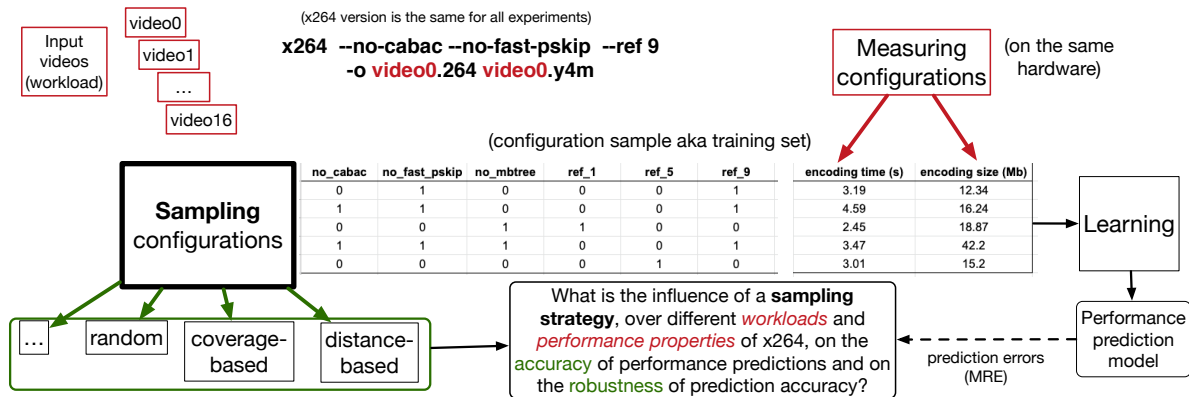


Fig. 1. Design study: sampling effect on performance predictions of x264 configurations

considered. Overall, a true uniform random sampling may be hard to realize, especially for large configurable systems. At the scale of the x264 study [24], though, uniform sampling is possible (the whole population is 1,152 configurations). The specific question we explore here is whether random is effective for learning (in case it is applicable as in x264).

When random sampling is not applicable, several alternate techniques have been proposed typically by sacrificing some uniformity for a substantial increase in performance.

Solver-based. Many works rely on off-the-shelf constraint solver, such as SAT4J [30] or Z3 [11], for sampling. For instance, a random seed can be set to the Z3 solver and internally influences the variable selection heuristics, which can have an effect on the exploration of valid configurations. Henard *et al.* noticed that solvers' internal order yields non-uniform (and predictable) exploration of the configuration space [17]. Hence, these strategies do not guarantee true randomness as in uniform random sampling. Often the sample set consists of a locally clustered set of configurations.

Randomized solver-based. To weaken the locality drawback of solver-based sampling, Henard *et al.* change the order of variables and constraints at each solver run. This strategy, called *randomized solver-based sampling* in [24], increases diversity of configurations. Though it cannot give any guarantees about randomness, the diversity may help to capture important interactions between options for performance prediction.

Coverage-based. sampling aims to optimize the sample with regards to a coverage criterion. Many criteria can be considered such as statement coverage that requires the analysis of the source code. In this paper and as in [24], we rely on *t*-wise sampling [10, 23, 29]. This sampling strategy selects configurations to cover all combinations of *t* selected options. For instance, pair-wise ($t=2$) sampling covers all pairwise combinations of options being selected. There are different methods to compute *t*-wise sampling. As in [24], we rely on the implementation of Siegmund *et al.* [52].

Distance-based. Kaltenecker *et al.* [24] propose distance-based sampling. The idea is to cover the configuration space by selecting

configurations according to a given probability distribution (typically a uniform distribution) and a distance metric. The underlying benefit is that distance-based sampling can better scale compared to an enumerative-based random sampling, while the generated samples are closed to those obtained with a uniform random.

Diversified distance-based. sampling is a variant of distance-based sampling [24]. The principle is to increase diversity of the sample set by iteratively adding configurations that contain the least frequently selected options. The intended benefit is to avoid missing the inclusion of some (important) options in the process.

2.3 Sampling Effect on x264

This paper aims to replicate the study of Kaltenecker *et al.* [24]. While Kaltenecker *et al.* analyse a wide variety of systems from different domains, we focus on the analysis of a single configurable system. Compared to that paper, our study enables a deeper analysis of sampling strategies over the possible influences of inputs and performance properties. Specifically, we analyze a set of seventeen input videos and two non-functional properties. As in [24], we collect results of 100 independent experiments to increase statistical confidence and external validity. We aim at exploring whether the results obtained in [24] may be generalized over different variations of a single configurable system. Interestingly, numerous papers have specifically considered the x264 configurable system for assessing their proposals [14, 20, 21, 42, 45, 50, 58, 62, 63], but a fixed performance property or input video is usually considered. Jamshidi *et al.* [20] explore the impacts of versions, workloads, and hardware of several configurable systems, including x264. However, the effect of sampling strategies was not considered. In [21], sampling strategies for transfer learning were investigated; we are not considering such scenarios in our study. Overall, given a configurable system like x264, practitioners face the problem of choosing the right techniques for "sampling, measuring, learning". Specifically, we aim to understand what sampling strategy to use and whether there exists a sampling to rule any configuration space of x264.

The use of different input videos obviously changes the raw and absolute performance values, but it can also change the overall

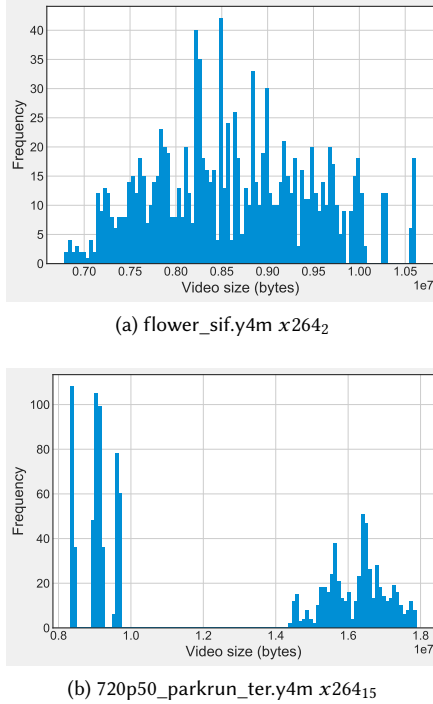


Fig. 2. The size distribution of 1,152 configurations of x264 over two input videos

distribution of configuration measurements. Figure 2 gives two distributions over two input videos for the performance property *size* for the whole population of valid configurations. Pearson correlation between the performance measurements of x264₂ and x264₁₅ is -0.35 while Spearman correlation is -0.021, suggesting a very weak correlation. The differences among distributions question the existence of a one-size-fits-all sampling capable of generating a representative sample set whatever the input videos or performance properties. Another hypothesis is that the way the sampling is done can pay off for some distributions but not for all.

Given the vast variety of possible input videos and performance properties that may be considered, the performance variability of configurations grows even more. Our aim is to investigate how such factors affect the overall prediction accuracy of different sampling strategies: *Is there a dominant sampling strategy for performance prediction of the same configurable system?*

3 DESIGN STUDY

In this section, we introduce our research questions, the considered subject system, and the experiment setup.

3.1 Research Questions

We conducted a series of experiments to evaluate six sampling strategies and to compare our results to the original results in [24]. We aim at answering the following two research questions:

- **(RQ1)** What is the influence of using different sampling strategies on the accuracy of performance predictions over different inputs and non-functional properties?

- **(RQ2)** What is the influence of randomness of using different sampling strategies on the robustness of prediction accuracy?

It is not new the assumption that the prediction accuracy of machine learning extensively depends on the sampling strategy. The originality of the research question is to what extent are performance prediction models of the *same* configurable system (here: x264) sensitive to other factors, such as different inputs and non-functional properties. To address RQ1, we analyze the sensitivity of the prediction accuracy of sampling strategies to these factors. Since most of the considered sampling strategies use randomness, which may considerably affect the prediction accuracy, RQ2 quantitatively compares whether the variances (over 100 runs) on prediction accuracy between different sampling strategies and sample sizes differ significantly. We show that the sampling prediction accuracy and robustness hardly depends on the definition of performance (*i.e.*, encoding time or encoding size). As in [24], we have excluded t-wise sampling from RQ2, as it is also deterministic in our setting and does not lead to variations.

3.2 Subject System

We conduct an in-depth study of x264, a popular and highly configurable video encoder implemented in C. We choose x264 instead of the other case studies documented in [24] because x264 demonstrated more promising accuracy results to the newest proposed sampling approach (*i.e.*, diversified distance-based sampling). With this study, we aim at investigating, for instance, whether diversified distance-based sampling also dominates across different variations of x264 (*i.e.*, inputs, performance properties). As benchmark, we encoded 17 different input videos from AVI to the H.264 codec and measured two quantitative properties (encoding time and encoding size).

- Encoding time (in short **time**): how many seconds x264 takes to encode a video.
- Encoding size of the output video (in short **size**): compression size (in bytes) of an output video in the H.264 format.

All measurements have been performed over the same version of x264 and on a grid computing infrastructure called IGRIDA¹. Importantly, we used the same hardware characteristics for all performance measurements. In Table 1, we provide an overview of the encoded input videos. This number of inputs allows us to draw conclusions about the practicality of sampling strategies for diversified conditions of x264.

To control measurement bias while measuring execution time, we have repeated the measurements several times. We report in Table 1 how many times *#times* we have repeated the time measurements of 1,152 configurations for each video input. Since measurements are costly, we set the repetition to up 30 times given 1-hour restriction from where we got the *#times* in Table 1. We repeated the measurements at least three times and at most 23 times and retained the average execution time for each configuration. The *stability* column reports how pairs of configuration measurements for a given video are correlated (Pearson correlation, on average). It allows us to control that measurements are indeed stable. For

¹<http://igrida.gforge.inria.fr/>

	video	#times	stability
x264 ₀	bridge_far_cif.y4m	5	0.997824
x264 ₁	ice_cif.y4m	5	0.968627
x264 ₂	flower_sif.y4m	5	0.971209
x264 ₃	claire_qcif.y4m	5	0.898657
x264 ₄	sintel_trailer_2k_480p24.y4m	9	0.995809
x264 ₅	football_cif.y4m	5	0.984586
x264 ₆	crowd_run_1080p50.y4m	3	0.999077
x264 ₇	blue_sky_1080p25.y4m	8	0.996295
x264 ₈	FourPeople_1280x720_60.y4m	11	0.995452
x264 ₉	sunflower_1080p25.y4m	4	0.998713
x264 ₁₀	deadline_cif.y4m	23	0.994600
x264 ₁₁	bridge_close_cif.y4m	5	0.997647
x264 ₁₂	husky_cif.y4m	5	0.984500
x264 ₁₃	tennis_sif.y4m	5	0.961498
x264 ₁₄	riverbed_1080p25.y4m	3	0.997521
x264 ₁₅	720p50_parkrun_ter.y4m	8	0.997049
x264 ₁₆	soccer_4cif.y4m	16	0.996276

Table 1. Overview of encoded input videos including the number of times we measured the encoding time in order to ensure we have a stable set of measurements (according to correlation results), independent of the machine.

example, measurements have been repeated 5 times for video 0 (bridge_far_cif.y4m) and the 5 measurements are highly correlated. Video 3 (claire_qcif.y4m) can be considered as an outlier (it may be due to the randomness of x264 given this specific input) because the correlation measure is below 0.9. However the correlation remains very high and close to 1. We provide the variability model and the measurements of each video input for encoding time and encoding size on our supplementary website.

3.3 Experiment Setup

In our experiments, the independent variables are the choice of the input videos, the predicted non-functional-property, the sample strategies and the sample sizes.

For comparison, we used the same experiment design them in Kaltenecker et al. [24]. To evaluate the accuracy of different sampling strategies over different inputs and non-functional-properties, we conducted experiments using three different sample sizes. To be able to use the same sample sizes for all sampling strategies, we consider the sizes from t -wise sampling with $t=1$, $t=2$, and $t=3$. As described in Section 2.2, t -wise sampling covers all combinations of t configuration options being selected. We learn a performance model based on the sample sets along with the corresponding performance measurements defined by the different sampling strategies. Although several machine-learning techniques have been proposed in the literature with this purpose [45], such as linear regression [9, 19, 20, 24–26, 32, 38, 49, 51, 53, 66], classification and regression trees (CART) [1, 15, 20, 27, 38–41, 47, 50, 55, 58, 59, 62, 63, 67–70], and random forest [2, 5, 48, 60, 62]. In this paper, we use step-wise multiple linear regression [51] as in [24]. According to Kaltenecker et al. [24], multiple linear regression is often as accurate as CART and random forests.

To calculate the prediction error rate, we use the resulting performance models to predict the performance of the entire dataset of valid configurations C , $t \subset C$. We calculate the error rate of a prediction model in terms of the *mean relative error* (MRE - Equation 1). MRE is used to estimate the accuracy between the exact measurements and the predicted one.

$$MRE = \frac{1}{|C|} \sum_{c \in C} \frac{|measured_c - predicted_c|}{measured_c} \quad (1)$$

Where C is the set of all valid configurations used as the validation set, and $measured_c$ and $predicted_c$ indicate the measured and predicted values of performance for configuration c with $c \in C$, respectively. The exact value of $measured_c$ is measured at runtime while running the configuration c , and the predicted values of $predicted_c$ is computed based on the model built with a sample of configurations t (see Section 2.2). To address *RQ1*, we computed the mean error rate for each input video and sample size. A lower error rate indicates a higher accuracy. Then, we use a Kruskal-Wallis test [28] and pair-wise one-sided Mann-Whitney U tests [33] to identify whether the error rate of two sampling strategies differs significantly ($p < 0.05$) [3]. In addition, we compute the effect size \hat{A}_{12} [64] (small(>0.56), medium(>0.64), and large(>0.71)) to easily compare the error rates of two sampling strategies.

To address *RQ2*, we compute the variance across the error rates over 100 runs. A lower variance indicates higher robustness. First, we use Levene’s test [31] to identify whether the variances of two sampling strategies differ significantly from each other. Then, for these sampling strategies, we perform a one-sided F-tests [54] to compare pair-wisely the variance between sampling strategies.

All sampling and learning experiments have been performed on the same machine with Intel Core i7 CPU 2,2 GHz and 4GB RAM. To reduce fluctuations in the values of dependent variables caused by randomness (e.g., the random generation of input samples), we evaluated each combination of the independent variables 100 times. That is, for each input video, non-functional property, sampling strategy and sampling size, we instantiated our experimental settings and measured the values of all dependent variables 100 times with random seeds from 1 to 100.

4 RESULTS

We compare six sampling strategies: t -wise, solver-based, randomized solver-based, distance-based, diversified distance-based, and random. Next, We present the results regarding prediction accuracy (*RQ1*, Section 4.1) and robustness (*RQ2*, Section 4.2).

4.1 Results RQ1—Prediction Accuracy

In Tables 2 and 4, we show the MRE for the different sampling strategies and sample sizes for both encoding time and encoding size, respectively. In the bottom row, we provide the MRE mean over all input videos. As in Kaltenecker et al. [24], for each sample-set size we highlight the lowest, statistically significant MRE in bold.

4.1.1 Input Sensitivity: Encoding Time. Random sampling performs best to all other sampling strategies or similar to the best one for $t=1$, $t=2$ and $t=3$ (except for a few input videos – x264₁, x264₃, x264₇, x264₈, x264₉, x264₁₀, x264₁₁ for $t=1$). We observe that for

Video	Coverage-based			Solver-based			Randomized solver-based			Distance-based			Diversified distance-based			Random		
	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$
x264 ₀	18.2%	13.9%	13.4%	24.0%	27.0%	27.5%	22.3%	19.9%	24.3%	16.5%	12.7%	10.6%	16.3%	8.8%	8.2%	16.7%	9.2%	8.2%
x264 ₁	15.4%	13.2%	12.1%	26.9%	23.7%	24.9%	21.4%	21.5%	23.2%	17.3%	14.2%	9.5%	17.4%	9.8%	8.7%	16.1%	9.2%	8.7%
x264 ₂	29.3%	10.3%	9.7%	21.4%	19.4%	16.4%	19.1%	19.6%	19.4%	17.4%	11.4%	9.8%	17.6%	9.6%	9.3%	15.3%	9.5%	9.3%
x264 ₃	21.4%	13.7%	10.1%	25.2%	25.3%	26.4%	16.4%	22.3%	24.8%	13.6%	10.7%	10.2%	12.8%	9.8%	9.7%	14.5%	9.8%	9.2%
x264 ₄	21.8%	12.3%	14.4%	23.9%	21.2%	22.0%	18.3%	21.1%	22.5%	14.2%	11.7%	9.7%	13.9%	10.1%	8.9%	13.9%	9.4%	8.8%
x264 ₅	26.1%	14.1%	13.2%	28.8%	23.2%	24.1%	21.8%	22.5%	23.3%	16.4%	13.4%	11.4%	16.8%	10.7%	9.5%	15.7%	10.0%	9.3%
x264 ₆	25.9%	18.1%	8.6%	23.6%	28.5%	29.1%	18.2%	21.6%	24.9%	13.7%	9.9%	9.0%	13.2%	8.8%	7.8%	12.6%	8.0%	7.3%
x264 ₇	23.3%	14.2%	12.0%	20.2%	25.3%	26.1%	15.3%	23.0%	23.8%	12.2%	9.2%	7.2%	10.8%	8.5%	7.2%	11.4%	8.2%	7.3%
x264 ₈	20.8%	13.1%	11.5%	20.3%	22.7%	23.6%	16.7%	23.4%	23.4%	12.6%	10.4%	9.6%	11.1%	9.3%	8.3%	12.0%	8.7%	7.6%
x264 ₉	23.4%	13.2%	5.6%	22.1%	28.6%	29.7%	16.8%	24.2%	25.3%	11.4%	6.5%	6.5%	9.2%	5.8%	5.4%	10.9%	6.6%	5.4%
x264 ₁₀	21.9%	12.3%	9.3%	22.6%	23.2%	24.0%	17.9%	22.4%	24.3%	14.0%	10.2%	9.7%	13.5%	9.4%	8.9%	14.0%	9.0%	8.8%
x264 ₁₁	21.1%	12.6%	10.3%	25.7%	23.5%	23.8%	20.0%	21.1%	24.7%	13.3%	10.8%	10.4%	13.0%	10.1%	9.7%	13.9%	9.4%	9.1%
x264 ₁₂	25.4%	13.4%	10.4%	26.2%	21.2%	21.6%	19.8%	20.6%	20.9%	16.2%	13.7%	10.9%	16.3%	11.4%	9.1%	15.0%	9.7%	8.5%
x264 ₁₃	16.4%	10.5%	10.0%	20.6%	18.8%	19.1%	18.3%	19.4%	19.8%	16.0%	13.9%	10.0%	16.2%	10.5%	9.6%	15.5%	9.7%	9.0%
x264 ₁₄	20.7%	16.9%	15.8%	34.3%	39.5%	40.6%	28.5%	29.7%	32.4%	18.1%	11.1%	9.6%	18.4%	7.8%	7.3%	17.4%	7.5%	7.2%
x264 ₁₅	26.2%	12.7%	11.1%	23.2%	26.5%	27.2%	20.3%	22.7%	25.1%	15.1%	11.9%	10.7%	14.8%	10.6%	9.5%	13.9%	9.1%	8.9%
x264 ₁₆	22.9%	12.3%	8.4%	22.1%	24.5%	25.2%	18.0%	22.2%	23.6%	13.4%	9.4%	8.9%	12.6%	8.5%	7.8%	12.5%	8.1%	7.4%
Mean	22.4%	13.3%	10.9%	24.2%	24.8%	25.4%	19.4%	22.2%	23.9%	14.8%	11.3%	9.6%	14.3%	9.4%	8.5%	14.2%	8.9%	8.2%

Table 2. Error rates of t -wise, (randomized) solver-based, (diversified) distance-based, and random sampling for the prediction of time for 17 input videos of x264. The bottom row contains the MRE mean across all input videos. The best results per input video and sample size are highlighted in bold if the Mann-Whitney U test reported a significant difference ($p < 0.05$).

Mann-Whitney U test [p value (\hat{A}_{12})]																		
	Coverage-based			Solver-based			Randomized solver-based			Distance-based			Diversified distance-based			Random		
	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$
Coverage-based																		
Solver-based						0 (0.99)			0 (1.00)			0 (0.98)			0 (1.00)			0 (0.65)
Randomized solver-based																		
Distance-based																		
Diversified distance-based																		
Random																		

Table 3. p values from a one-sided pair-wise Mann-Whitney U test for the property encoding time, where we tested pair-wisely whether the error rate of the sampling strategy from the row is significantly lower than the error rate of the sampling strategy from the column, for different sample sizes. The effect size is included for every significant result ($p < 0.05$), where we consider the effect as small, medium, and large when \hat{A}_{12} is over 0.56, 0.64, and 0.71, respectively.

$t=1$ diversified distance-based sampling outperforms random sampling for seven input videos (x264₀, x264₃, x264₇, x264₈, x264₉, x264₁₀, x264₁₁). Overall, diversified distance-based sampling produces partially good results (close to random sampling). Diversified distance-based sampling outperforms the pure distance based sampling, however their error rates are very similar for $t=1$.

Solver-based sampling results in inaccurate performance models for all input videos and sample-set sizes. t -wise sampling performs overall better than solver-based sampling; randomized solver-based sampling performs best when only a very limited number of samples are considered (i.e., $t=1$).

Table 3 reports the p value (\hat{A}_{12}) of the Mann-Whitney U test. This table shows whether the sampling strategy of the row has a significantly lower error rate than the sampling strategy of the column. To this end, we first performed Kruskal-Wallis tests for all sample sizes ($t=1$, $t=2$, and $t=3$). Then, whether we identified p values less than 0.05, indicating that, at least, two strategies differ significantly for each sample size [24], we performed one-sided

Mann-Whitney U tests pair-wisely and, if significant ($p < 0.05$), we report the effect sizes in Table 3.

The first row shows that t -wise sampling has significantly lower error rates than solver-based sampling and randomized solver-based sampling for $t=2$ and $t=3$, with large effect sizes. In the second and third rows, we observe that solver-based sampling performs significantly better than randomized solver-based sampling for $t=3$ with medium effect sizes; and randomized solver-based leads to lower error rates than t -wise sampling for $t=1$, with a large effect size. In the last three rows, we see that (diversified) distance-based sampling and random sampling lead to lower error rates than t -wise sampling, solver-based sampling, and randomized solver-based sampling for $t=1$, $t=2$ and $t=3$, with large effect sizes.

The pure distance-based sampling leads to higher error rates than diversified distance-based sampling and random sampling, for all sample sizes. However, comparing diversified distance-based sampling to random sampling, we see in the last row that random sampling has significantly lower error rates with small effect sizes for all sample sizes. The small effect sizes indicate that diversified

Video	Coverage-based			Solver-based			Randomized solver-based			Distance-based			Diversified distance-based			Random		
	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$
x2640	12.3	11.6	11.1	12.3	11.4	11.3	25.1	12.7	13.3	25.3	12.5	10.6	23.3	10.6	9.2	13.1	9.8	9.1
x2641	4.0	3.9	3.8	3.1	3.8	3.8	1.7	3.8	3.8	4.0	4.0	3.8	3.9	3.8	3.8	3.9	3.8	3.8
x2642	14.9	14.3	4.8	5.1	4.7	4.7	15.9	4.7	4.6	14.3	14.0	10.2	13.8	12.0	4.7	7.6	4.7	4.6
x2643	8.6	8.3	7.8	8.1	7.3	7.4	11.2	7.6	7.4	9.9	9.3	8.0	9.6	8.3	7.5	7.7	7.4	7.3
x2644	18.4	16.7	6.6	4.5	6.8	6.8	14.1	6.7	6.7	17.5	16.7	7.0	16.9	6.9	6.9	7.8	6.9	6.9
x2645	11.3	11.0	10.8	4.9	6.6	5.7	12.3	9.4	4.8	11.8	11.5	10.9	11.6	10.6	10.0	9.4	6.4	5.2
x2646	24.6	5.3	5.2	5.4	5.4	5.3	25.6	5.3	5.3	17.6	16.8	5.5	16.1	5.4	5.4	6.3	5.3	5.3
x2647	9.4	9.0	8.7	8.1	8.4	8.3	8.4	8.2	8.2	9.4	9.4	8.9	9.3	8.6	8.5	9.1	8.4	8.3
x2648	10.4	9.7	8.9	8.7	8.0	8.1	11.2	7.6	8.0	12.4	12.0	9.5	12.0	9.9	8.5	8.5	8.3	8.2
x2649	11.6	10.5	9.5	7.6	8.6	8.5	6.9	8.4	8.4	11.3	11.6	9.6	10.8	9.7	8.7	8.8	8.5	8.4
x26410	5.2	5.2	4.9	5.2	5.0	4.8	5.0	4.6	4.6	6.0	5.8	5.0	5.7	5.1	4.7	4.9	4.6	4.6
x26411	12.4	11.8	11.1	11.1	10.8	11.0	8.8	9.9	11.4	12.8	11.8	9.0	12.0	10.2	8.6	10.9	9.4	8.8
x26412	25.7	3.6	3.6	5.3	3.5	3.6	28.9	3.6	3.5	16.5	14.6	3.5	15.4	3.5	3.4	4.8	3.5	3.4
x26413	4.7	4.7	4.6	4.5	4.7	4.7	5.4	4.8	4.7	5.1	5.0	4.8	5.0	4.7	4.7	5.0	4.7	4.6
x26414	10.2	9.6	9.4	5.1	7.4	8.8	3.6	9.6	9.5	10.6	10.6	10.0	9.8	9.6	9.6	9.3	9.0	9.5
x26415	4.1	4.0	4.0	7.5	4.5	4.3	40.9	4.3	4.2	21.7	8.3	4.1	19.1	4.1	4.1	5.4	4.2	4.1
x26416	8.3	8.1	7.9	7.7	7.8	7.6	9.2	7.7	7.6	8.8	8.7	8.2	8.7	7.9	7.7	8.3	7.7	7.6
Mean	11.5	8.7	7.2	6.7	6.8	6.7	13.8	7.0	6.8	12.6	10.7	7.6	12.0	7.7	6.8	7.7	6.6	6.5

Table 4. Error rates of t -wise, (randomized) solver-based, (diversified) distance-based, and random sampling for the prediction of `size` for 17 input videos of x264. The bottom row contains the MRE mean across all input videos. The best results per input video and sample size are highlighted in bold if the Mann-Whitney U test reported a significant difference ($p < 0.05$).

	Mann-Whitney U test [p value (\hat{A}_{12})]																	
	Coverage-based			Solver-based			Randomized solver-based			Distance-based			Diversified distance-based			Random		
	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$
Coverage-based							0.01 (0.52)			10 ⁻³³ (0.62)	10 ⁻¹⁷⁶ (0.78)	10 ⁻⁷⁸ (0.68)	10 ⁻⁰⁸ (0.55)					
Solver-based	10 ⁻²³² (0.82)	10 ⁻²⁶⁸ (0.85)	10 ⁻⁴³ (0.64)				10 ⁻¹⁵² (0.76)			10 ⁻²⁴⁷ (0.83)	10 ⁻¹⁴⁵ (0.92)	10 ⁻¹⁴⁵ (0.75)	10 ⁻²¹⁹ (0.81)	10 ⁻⁸⁶ (0.69)				10 ⁻⁹⁵ (0.70)
Randomized solver-based		10 ⁻²⁷² (0.85)	10 ⁻⁶¹ (0.66)			10 ⁻⁰⁵ (0.54)				10 ⁻¹³ (0.57)	10 ⁻¹³ (0.93)	10 ⁻¹⁸⁴ (0.79)	10 ⁻¹⁸⁴ (0.52)	10 ⁻⁹⁴ (0.70)	10 ⁻⁰⁶ (0.54)			
Distance-based																		
Diversified distance-based		10 ⁻⁷⁶ (0.68)	10 ⁻⁷⁰ (0.68)							10 ⁻¹⁰ (0.56)	10 ⁻²⁹³ (0.86)	10 ⁻¹⁹¹ (0.79)						
Random	10 ⁻¹⁸⁰ (0.78)	0 (0.92)	10 ⁻¹⁸³ (0.79)		10 ⁻²⁸ (0.61)	10 ⁻⁵⁸ (0.66)	10 ⁻⁸³ (0.69)	10 ⁻²² (0.60)	10 ⁻³⁷ (0.63)	10 ⁻²²⁸ (0.82)	0 (0.96)	0 (0.88)	10 ⁻¹⁷⁹ (0.78)	10 ⁻¹⁹³ (0.79)	10 ⁻⁷¹ (0.68)			

Table 5. p values from a one-sided pair-wise Mann-Whitney U test for the property encoding `size`, where we tested pair-wisely whether the error rate of the sampling strategy from the row is significantly lower than the error rate of the sampling strategy from the column, for different sample sizes. The effect size is included for every significant result ($p < 0.05$), where we consider the effect as small, medium, and large when \hat{A}_{12} is over 0.56, 0.64, and 0.71, respectively.

distance-based sampling can reach nearly the same error rates as random sampling.

For the property encoding time, uniform random sampling yields the most accurate performance models. Diversified distance-based sampling produces good results when a very limited number of samples are considered (*i.e.*, $t=1$) and almost reaches the accuracy of random when the sample sizes increase.

4.1.2 Input Sensitivity: Encoding Size. Random sampling and randomized solver-based sampling perform best to all other sampling strategies or similar to the best sampling for all input videos for sample sizes $t=2$ and $t=3$ (except for $x254_{14}$). For these sample sizes, most of the error rates of solver-based and diversified distance-based come close to random and randomized solver-based (*e.g.*, $x264_1$ and $x264_6$). For $t=1$, solver-based sampling leads to lower error rates for most of the inputs.

It is important to notice that the results in Table 4 are quite unstable compared to the results in Table 2. For example, random is the

best sampling strategy for the input video $x264_{12}$, while a randomized solver-based sampling win for the input video $x264_1$, and t -wise sampling win for the input video $x264_{15}$. Also, the accuracy heavily depends on the sampling size. For example, for the input video $x264_{14}$, for $t=1$ randomized solver-based sampling yields the most accurate performance models; while for $t=2$ and $t=3$ solver-based outperforms the other sampling strategies.

In Table 5, we apply one-sided Mann-Whitney U tests pair-wisely to compare pairs of sampling strategies. The first row shows that t -wise sampling has a significantly lower error rate than distance-based sampling for $t=1$, $t=2$, and $t=3$ with small, large and medium effect sizes, respectively; t -wise sampling has also significantly lower error rates than diversified distance-based for $t=1$.

In the second row, we see that solver-based sampling leads to lower error rates than all other sampling strategies for $t=1$ with medium and large effect sizes. Solver-based, randomized solver-based, and diversified distance-based lead to lower error rates than t -wise for $t=2$ and $t=3$; and distance-based for $t=1$, $t=2$ and $t=3$. Overall, distance-based sampling leads to higher or similar error rates than other sampling strategies. Finally, we can notice that

	F-test (p value)														
	Solver-based			Randomized solver-based			Distance-based			Diversified distance-based			Random		
	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$
Solver-based															
Randomized solver-based	10^{-252}	10^{-125}													
Distance-based	10^{-238}	10^{-272}			10^{-38}										
Diversified distance-based	10^{-259}				10^{-238}	10^{-89}	0.05	10^{-100}	10^{-179}						
Random	10^{-170}				10^{-195}	10^{-46}		10^{-71}	10^{-118}						

Table 6. p values from a one-sided pair-wise F-test for **time**: we tested pair-wisely whether the variances of the error rate of the sampling from the row is significantly lower than the one from the column, for different sample sizes.

	F-test (p value)														
	Solver-based			Randomized solver-based			Distance-based			Diversified distance-based			Random		
	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$
Solver-based															
Randomized solver-based	10^{-38}	10^{-69}	10^{-40}					10^{-94}							
Distance-based	10^{-75}			10^{-09}											
Diversified distance-based	10^{-89}	10^{-27}	10^{-136}	10^{-14}		10^{-34}	0.04	10^{-202}	10^{-144}						
Random	10^{-126}	10^{-75}	10^{-252}	10^{-31}		10^{-107}	10^{-09}	10^{-302}	10^{-263}	10^{-05}	10^{-15}	10^{-25}			

Table 7. p values from a one-sided pair-wise F-test for **size**: we tested pair-wisely whether the variances of the error rate of the sampling from the row is significantly lower than the one from the column, for different sample sizes.

on average random leads to significantly lower error rates than all other strategies (except for solver-based sampling for $t=1$).

For the property encoding size, random sampling and randomized solver-based sampling outperform all other sampling strategies for most of the input videos with sample sizes $t=2$ and $t=3$; and solver-based sampling outperforms for sample sizes $t=1$. Overall, random, randomized solver-based, solver-based, and diversified distance-based present good and similar accuracy for $t=2$ and $t=3$. Differently from our previous results (for **time**), there is not a clear winner.

4.2 Results RQ2—Robustness

We repeated each experiment 100 times for each sampling strategy and sample size, from which we obtained the distribution of mean error rates. We use the Levene’s test to check whether there are significantly different variances between sampling strategies. In cases where we found significant variances, we performed pair-wisely one-sided F-tests. We show the results in Tables 6 and 7.

4.2.1 Encoding Time. In the second and third rows of Table 6, we can see that randomized solver-based sampling and distance-based sampling have a significantly lower variance than solver-based sampling for $t=1$ and $t=2$; the same also applies for random with $t=1$. In the two last rows, we observe that diversified distance-based and random have a lower variance than solver-based for $t=1$; both sampling have also lower variance than randomized solver-based and distance-based, for $t=2$ and $t=3$. Diversified distance-based has a significantly lower variance than random sampling for all sample sizes. Finally, random has a lower variance than t -wise for $t=1$; and randomized solver-based and distance-based for $t=2$ and $t=3$.

For the property encoding time, diversified distance-based sampling is more robust than the other sampling strategies.

4.2.2 Encoding Size. In the first row of Table 7, we can see that solver-based sampling has a significantly lower variance than distance-based sampling for $t=2$. In the second row, we can see that randomized solver-based sampling has a significantly lower variance than solver-based sampling on all sample sizes. Randomized solver-based sampling has also a significantly lower variance than distance-based sampling for $t=2$ and $t=3$; and diversified distance-based sampling for $t=2$. When it comes to the diversified distance-based sampling, it leads to a significantly lower variance than the other sampling strategies, except for randomized solver-based ($t=2$) and random. In the last row, we observe that random has the lowest variance.

For the property encoding size, uniform random sampling is more robust than the other sampling strategies.

5 DISCUSSION

The surprising effectiveness of coverage-based sampling. Why t -wise sampling is more effective for the input video $x264_{15}$? For $x264_{15}$, the performance distribution is rather unique compared to the others (see Figure 2b, page 4). To investigate why t -wise is more effective on this distribution, we analyzed the set of selected features by t -wise. We aim at understanding which features are frequently included in the sampling and how its frequencies differ from random. We recall that t -wise sampling is not subject to randomness and is the same for any video and for encoding time or size. We observed that the t -wise sampling prioritizes the feature `no_mbtrees` to false, *i.e.*, `no_mbtrees` is most of the time deactivated (it is not the

Video	Coverage-based			Solver-based			Randomized solver-based			Distance-based			Diversified distance-based			Random		
	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$
<i>sintel_trailer(x264₄)</i>	21.8 %	12.3 %	14.4 %	23.9 %	21.2 %	22.0 %	18.3 %	21.1 %	22.5 %	14.2 %	11.7 %	9.7 %	13.9 %	10.1 %	8.9 %	13.9 %	9.4 %	8.8 %
<i>sintel_trailer(x264[24])</i>	20.9 %	11.9 %	10.9 %	26.2 %	40.4 %	42.2 %	18.5 %	22.2 %	33.2 %	14.7 %	10.0 %	9.4 %	12.6 %	8.8 %	9.0 %	13.5 %	9.2 %	9.1 %

Table 8. Error rates of t -wise, (randomized) solver-based, (diversified) distance-based, and random sampling for the prediction of encoding time for the Sintel trailer input video (734 MB). The best results are highlighted in bold if the Mann-Whitney U test reported a significant difference ($p < 0.05$).

case in random that has a good balance). To better understand the importance of `no_mbtrees` for `x26415`, we use (1) a random forest for computing so-called feature importance [6, 12, 36, 44] and (2) polynomial regressions for computing coefficients of options and pair-wise interactions among options [13, 49, 67]. We found that `no_mbtrees` has a feature importance of 0.97 (out of 1) for encoding size (see our supplementary website). We also found that the feature `no_mbtrees` has a strong negative effect on video `x26415`. As for solver-based sampling, t -wise sampling relies on clustered zones (see Section 2.2) that luckily cover the most influential options. The clustering zones must contain features with high significance on performance which may explain high accuracy results. Specifically, the sample-set of t -wise is locally clustered around the configurations with `no_mbtrees` to false, which is a very effective strategy for `x26415`. However, for other input videos, the coverage criterion is not well-suited.

The surprising effectiveness of solver-based sampling. For the property encoding size, we observe good accuracy predictions for three sampling strategies: random, randomized solver-based, and solver-based. Why solver-based sampling strategies are sometimes more effective than random for size? We further analyzed the set of frequently selected features by these strategies for both experiments of time and size. We observe that the randomized solver-based and solver-based strategies were very effective to capture “clustered” effects of important options for size. In particular, randomized solver-based tend to follow a similar strategy as t -wise: `no_mbtrees` option is most of the time false in the sample set. For $t=1$, `no_mbtrees` option is even always set to false. The employed solver-based sampling strategy relies on a random seed which defines clustering zones (see Section 2.2) that luckily cover the most influential options. We did observe strong deviations of options’ frequencies in the samples of solver-based sampling strategy compared to random (see our supplementary website). For the prediction of encoding time, the clustering mostly contain features with low significance which may explain the worst accuracy results.

There are several solvers available and each solver has its own particularities, *i.e.*, the way they pick and thus cluster configurations can drastically change. The assumption that a certain sampling is the best one by comparing it to a solver-based sampling is not reliable since solver-based strategies may cluster by accident the (un-)important options to a specific performance property. Our study calls to propose sampling strategies with predictable and documented properties of the frequencies of options in the sample (*e.g.*, uniform random sampling or distance-based sampling).

Is there a “dominant” sampling strategy when the system performance is measured over the variation of hardware and the version of the target system? We compare our results with

the results obtained in Kaltenecker et al. [24] to encode the *same* (`x2644`) input video Sintel trailer (734 MB). Their measurements were performed on different hardware (Intel Core Q6600 with 4 GB RAM (Ubuntu 14.04)) and over a different version of `x264`. Diversity distance-based sampling yields the best performance in their settings. In our case, random sampling outperforms diversity distance-based sampling (see Table 8). Thus, we can hypothesize that *sampling strategies are also sensitive to different hardware characteristics and software versions (in addition to targeted performance properties and input videos)*. Further experiments on different hardware and versions are needed though to confirm this hypothesis. There are several questions that arise from this: To what extent are sampling strategies sensitive to different hardware characteristics and software versions? Are the interactions and influence of configuration options on performance consistent across different software versions? Do we experience the same for different performance properties?

Answering RQ1 (accuracy) Is there a “dominant” sampling strategy for the `x264` configurable system whatever inputs and targeted quantitative properties? For the property encoding time, there is a dominant sampling strategy (*i.e.*, random sampling), and thus the sampling can be reused whatever the input video is. For the property encoding size, although the results are similar around some sampling strategies, they differ in a noticeable way from encoding time and suggest a higher input sensitivity. Overall, random is the state-of-the-art sampling strategy but is not a dominant sampling strategy in all cases, *i.e.*, the ranking of dominance changes significantly given different inputs, properties and sample sizes. A possible hypothesis is that individual options and their interactions can be more or less influential depending on input videos, thus explaining the variations’ effect of sampling over the accuracy. Our results pose a new challenge for researchers: Identifying what sampling strategy is the best given the possible factors influencing the configurations’ performances of a system.

Answering RQ2 (robustness). We have quantitatively analyzed the effect of a sampling strategy over the prediction variance. Overall, random (for size) and diversified distance-based (for time and size) have higher robustness. We make the observation that uniform random sampling is not necessarily the best choice when robustness should be considered (but it is for accuracy). In practical terms, practitioners may have to find the right balance between the two objectives. As a sweet-spot between accuracy and robustness, diversified distance-based sampling (for time), and either random or randomized solver-based sampling (for size) are the best candidates. We miss however an actionable metric that could take both accuracy and robustness into account.

Our recommendations for practitioners are as follows:

- uniform random sampling is a very strong baseline whatever the inputs and performance properties. In the absence of specific-knowledge, practitioners should rely on this dominant strategy for reaching high accuracy;
- in case uniform random sampling is computationally infeasible, distance-based sampling strategies are interesting alternatives;
- the use of other sampling strategies does not pay off in terms of prediction accuracy. When robustness is considered as important, uniform random sampling is not the best choice and here we recommend diversified distance-based sampling.

The impacts of our results on configuration and performance engineering research are as follows:

- as uniform random sampling is effective for learning performance prediction models, additional research effort is worth doing to make it scalable for large instances. Recent results [18, 37, 43, 46] show some improvements, but the question is still open for very large systems (*e.g.*, Linux);
- some sampling strategies are surprisingly effective for specific inputs and performance properties. Our insights suggest the existence of specific sampling strategies that could prioritize specific important (interactions between) options. An open issue is to discover them for any input or performance property;
- performance measurements with similar distributions may be grouped together to enable the search for dominant sampling strategies;
- beating random is possible but highly challenging in all situations;
- it is unclear how factors such as the version or the hardware influence the sensitivity of the sampling effectiveness (and how such influence differs from inputs and performance properties);
- we warn researchers that the effectiveness of sampling strategies for a given configurable system can be biased by the workload and the performance property used.

6 THREATS TO VALIDITY

Despite the effort spent on the replication of the experiments in [24], we describe below some *internal* and *external* threats to the validity of this study.

Internal Validity. A threat to the internal validity of this study is the selection of the learning algorithm and its parameter settings which may affect the performance of the sampling strategy. We used the state-of-the-art step-wise multiple linear regression algorithm from [51] which has shown promising results in this field [24, 45]. However, we acknowledge that the use of another algorithm may lead to different results. Still, conducting experiments with other algorithms and tuning parameters is an important next step, which is part of our future work.

As another internal threat to validity, our results can be subject to measurement bias. While the property encoding size is stable, the measurement of the property encoding time deserves careful attention. To mitigate any measurement bias, we measured the time from all different input videos multiple times and used the

average in our learning procedure. Also, we control external factors like the hardware and workload of the machine by using a grid computing infrastructure called IGRIDA². We take care of using the same hardware and mitigate network-related factors. Using a public grid instead of a private cloud allows us to have control of the resources we used for measuring time. To additionally quantify the influence of measurement bias on our results, we also analyzed their Pearson correlation coefficients to prove we have a stable set of measurements.

To assess the accuracy of the sampling strategies and the effect of different input videos, we used MRE since most of the state-of-the-art works use this metric [45]. As in [24], we compute MRE based on the whole population of all valid configurations, *i.e.*, we include the sample t (used as training) to the testing set. However, according to several authors [15, 39, 40, 45][50, 57–59] the prediction error rate on new data is the most important measurement, *i.e.*, the error rate should not be computed over the training set t used to build the model. As future work, we plan to explore the effectiveness between four well-established resampling methods [45]: hold-out, cross-validation, bootstrapping, and dynamic sector validation.

External Validity. A threat to external validity is related to the used case study and the discussion of the results. Because we rely on a specific system and two performance properties, the results may be subject to this system and properties. However, we focused on a single system to be able at making robust and reliable statements about whether a specific sampling strategy can be used for different inputs and performance measurements. We conducted a discussion with all authors of the paper to avoid any confusing interpretation or misunderstanding of the results. Once we are able to demonstrate evidences to x264, we can then perform such an analysis also over other systems and properties to generalize our findings.

We set our dataset to up seventeen input videos and two properties due to budget restrictions. As with all studies, there is an inherent risk to not generalize our findings for other input videos and performance properties (*e.g.*, energy consumption). While the seventeen videos cover a wide range of different input particularities and provided consistent results across the experiments, this is a preliminary study in this direction and future work should consider additional inputs based on an in-depth qualitative study of video characteristics.

7 CONCLUSION

Finding a small sample that represents the important characteristics of the set of all valid configurations of a configurable system is a challenging task. Numerous sampling strategies have been proposed [45] and a recent study [24] have designed an experiment to compare six sampling strategies. However, this study did not investigate whether the strategies developed so far generalize across different inputs and performance properties of the same configurable system. To this end, we replicated this study to investigate the individual effects of 17 input videos on the prediction accuracy of two performance properties. Our results demonstrated that uniform random sampling dominates for most input videos and both performance properties, time and size. There are some cases for which random

²<http://igrida.gforge.inria.fr/>

sampling can be beaten with specific sampling. However, for the other sampling strategies, the prediction accuracy (*i.e.*, the ranking of sampling strategies) can dramatically change based on the input video and sampling size. It has practical implications since users of configurable systems feed different inputs and deal with different definitions of performance. Thus, we warn researchers about the sensitivity of a sampling strategy over workloads and performance properties of a configurable system. This work provides a new view of random sampling based on his unquestionable dominance across different inputs and performance properties. Distance-based sampling strategies are other relevant alternatives. Our replication of the original experiment design is a promising starting point for future studies of other configurable systems, consideration of other influential factors (versions, hardware), and the use of specific (*e.g.*, white-box) sampling strategies.

ACKNOWLEDGMENTS

This research was funded by the ANR-17-CE25-0010-01 VaryVary project. We thank Arnaud Blouin, Luc Lesoil, and Gilles Perrouin for their comments on an early draft of this paper.

REFERENCES

- [1] Mathieu Acher, Paul Temple, Jean-Marc Jezequel, José A Galindo, Jabier Martinez, and Tewfik Ziadi. 2018. VaryLaTeX: Learning Paper Variants That Meet Constraints. In *Proceedings of the 12th International Workshop on Variability Modelling of Software-Intensive Systems*. ACM, 83–88.
- [2] Benoit Amand, Maxime Cordy, Patrick Heymans, Mathieu Acher, Paul Temple, and Jean-Marc Jézéquel. 2019. Towards Learning-Aided Configuration in 3D Printing: Feasibility Study and Application to Defect Prediction. In *Proceedings of the 13th International Workshop on Variability Modelling of Software-Intensive Systems*. ACM, 7.
- [3] Andrea Arcuri and Lionel Briand. 2011. A Practical Guide for Using Statistical Tests to Assess Randomized Algorithms in Software Engineering. In *Proceedings of the 33rd International Conference on Software Engineering (ICSE '11)*. ACM, New York, NY, USA, 1–10.
- [4] Andrea Arcuri and Lionel Briand. 2012. Formal Analysis of the Probability of Interaction Fault Detection Using Random Testing. *IEEE Transactions on Software Engineering* 38, 5 (Sept 2012), 1088–1099.
- [5] Liang Bao, Xin Liu, Ziheng Xu, and Baoyin Fang. 2018. AutoConfig: automatic configuration tuning for distributed message systems. ACM, 29–40.
- [6] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [7] Supratik Chakraborty, Daniel J. Fremont, Kuldeep S. Meel, Sanjit A. Seshia, and Moshe Y. Vardi. 2015. On Parallel Scalable Uniform SAT Witness Generation. In *Tools and Algorithms for the Construction and Analysis of Systems TACAS'15 2015, London, UK, April 11–18, 2015*. Proceedings. 304–319.
- [8] Supratik Chakraborty, Kuldeep S Meel, and Moshe Y Vardi. 2013. A scalable and nearly uniform generator of SAT witnesses. In *International Conference on Computer Aided Verification*. Springer, 608–623.
- [9] Shiping Chen, Yan Liu, Ian Gorton, and Anna Liu. 2005. Performance prediction of component-based applications. *Journal of Systems and Software* 74, 1 (2005), 35–43.
- [10] Myra B Cohen, Matthew B Dwyer, and Shi, Jiangfan. 2008. Constructing Interaction Test Suites for Highly-Configurable Systems in the Presence of Constraints: A Greedy Approach. *IEEE TSE* 34, 5 (2008), 633–650.
- [11] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.
- [12] Aaron Fisher, Cynthia Rudin, and Francesca Dominici. 2018. All Models are Wrong, but Many are Useful: Learning a Variable's Importance by Studying an Entire Class of Prediction Models Simultaneously. arXiv:arXiv:1801.01489
- [13] Alexander Grebhahn, Carmen Rodrigo, Norbert Siegmund, Francisco J Gaspar, and Sven Apel. 2017. Performance-influence models of multigrid methods: A case study on triangular grids. *Concurrency and Computation: Practice and Experience* 29, 17 (2017), e4057.
- [14] Jianmei Guo, Krzysztof Czarnecki, Sven Apel, Norbert Siegmund, and Andrzej Wasowski. 2013. Variability-aware performance prediction: A statistical learning approach. In *Int' Conference on Automated Software Engineering (ASE)*. IEEE, 301–311.
- [15] Jianmei Guo, Jia Hui Liang, Kai Shi, Dingyu Yang, Jingsong Zhang, Krzysztof Czarnecki, Vijay Ganesh, and Huiqun Yu. 2017. SMTBEA: a hybrid multi-objective optimization algorithm for configuring large constrained software product lines. *Software & Systems Modeling* (22 Jul 2017). <https://doi.org/10.1007/s10270-017-0610-0>
- [16] Axel Halin, Alexandre Nuttinck, Mathieu Acher, Xavier Devroey, Gilles Perrouin, and Benoit Baudry. 2018. Test them all, is it worth it? Assessing configuration sampling on the JHipster Web development stack. *Empirical Software Engineering* (17 Jul 2018). <https://doi.org/10.1007/s10664-018-9635-4>
- [17] Christopher Henard, Mike Papadakis, Gilles Perrouin, Jacques Klein, Patrick Heymans, and Yves Le Traon. 2014. Bypassing the Combinatorial Explosion: Using Similarity to Generate and Prioritize T-Wise Test Configurations for Software Product Lines. *IEEE Trans. Software Eng.* (2014).
- [18] Ruben Heradio, David Fernández-Amorós, Christoph Mayr-Dorn, and Alexander Egyed. 2019. Supporting the statistical analysis of variability models. In *Proceedings of the 41st International Conference on Software Engineering, ICSE 2019, Montreal, QC, Canada, May 25–31, 2019*. 843–853.
- [19] Pooyan Jamshidi, Javier Cámara, Bradley Schmerl, Christian Kästner, and David Garlan. 2019. Machine Learning Meets Quantitative Planning: Enabling Self-Adaptation in Autonomous Robots. arXiv preprint arXiv:1903.03920 (2019).
- [20] Pooyan Jamshidi, Norbert Siegmund, Miguel Velez, Christian Kästner, Akshay Patel, and Yuvraj Agarwal. 2017. Transfer learning for performance modeling of configurable systems: an exploratory analysis. IEEE Press, 497–508. <http://dl.acm.org/citation.cfm?id=3155625>
- [21] Pooyan Jamshidi, Miguel Velez, Christian Kästner, and Norbert Siegmund. 2018. Learning to sample: exploiting similarities across environments to learn performance models for configurable systems. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ACM, 71–82.
- [22] Pooyan Jamshidi, Miguel Velez, Christian Kästner, Norbert Siegmund, and Prasad Kawthekar. 2017. Transfer Learning for Improving Model Predictions in Highly Configurable Software. In *12th IEEE/ACM International Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS@ICSE 2017, Buenos Aires, Argentina, May 22–23, 2017*. 31–41.
- [23] Martin Fagereng Johansen, Øystein Haugen, and Franck Fleurey. 2012. An algorithm for generating t-wise covering arrays from large feature models. In *Proceedings of the 16th International Software Product Line Conference on - SPLC '12 - volume 1, Vol. 1*. ACM, 46.
- [24] Christian Kaltenecker, Alexander Grebhahn, Norbert Siegmund, Jianmei Guo, and Sven Apel. 2019. Distance-Based Sampling of Software Configuration Spaces. In *Proceedings of the IEEE/ACM International Conference on Software Engineering (ICSE)*. ACM.
- [25] Sergiy Kolesnikov, Norbert Siegmund, Christian Kästner, and Sven Apel. 2017. On the relation of external and internal feature interactions: A case study. arXiv preprint arXiv:1712.07440 (2017).
- [26] Sergiy Kolesnikov, Norbert Siegmund, Christian Kästner, Alexander Grebhahn, and Sven Apel. 2019. Tradeoffs in modeling performance of highly configurable software systems. *Software & Systems Modeling* 18, 3 (01 Jun 2019), 2265–2283. <https://doi.org/10.1007/s10270-018-0662-9>
- [27] Thomas Krismayer, Rick Rabiser, and Paul Grünbacher. 2017. Mining constraints for event-based monitoring in systems of systems. IEEE Press, 826–831.
- [28] William H Kruskal and W Allen Wallis. 1952. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association* 47, 260 (1952), 583–621.
- [29] D.R. Kuhn, D.R. Wallace, and A.M. Gallo. 2004. Software fault interactions and implications for software testing. *IEEE Transactions on Software Engineering* 30, 6 (jun 2004), 418–421.
- [30] Daniel Le Berre and Anne Parrain. 2010. The SAT4J library, Release 2.2, System Description. *Journal on Satisfiability, Boolean Modeling and Computation* 7 (2010), 59–64. <https://hal.archives-ouvertes.fr/hal-00868136>
- [31] Howard Levene. 1961. Robust tests for equality of variances. *Contributions to probability and statistics. Essays in honor of Harold Hotelling* (1961), 279–292.
- [32] Max Lillack, Johannes Müller, and Ulrich W Eisenecker. 2013. Improved prediction of non-functional properties in software product lines with domain context. *Software Engineering 2013* (2013).
- [33] Henry B Mann and Donald R Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics* (1947), 50–60.
- [34] Flávio Medeiros, Christian Kästner, Márcio Ribeiro, Rohit Gheyi, and Sven Apel. 2016. A comparison of 10 sampling algorithms for configurable systems. In *Proceedings of the 38th International Conference on Software Engineering - ICSE '16*. ACM Press, Austin, Texas, USA, 643–654.
- [35] Marcilio Mendonca, Andrzej Wasowski, Krzysztof Czarnecki, and Donald Cowan. 2008. Efficient compilation techniques for large scale feature models. In *Proceedings of the 7th international conference on Generative programming and component engineering*. ACM, 13–22.

- [36] Christoph Molnar. 2019. *Interpretable Machine Learning*. <https://christophm.github.io/interpretable-ml-book/>.
- [37] Daniel-Jesus Munoz, Jeho Oh, Mónica Pinto, Lidia Fuentes, and Don S. Batory. 2019. Uniform random sampling product configurations of feature models that have numerical features. In *Proceedings of the 23rd International Systems and Software Product Line Conference, SPLC 2019, Volume A, Paris, France, September 9-13, 2019*. 39:1–39:13. <https://doi.org/10.1145/3336294.3336297>
- [38] I Made Murwantara, Behzad Bordbar, and Leandro L. Minku. 2014. Measuring Energy Consumption for Web Service Product Configuration. In *Proceedings of the 16th International Conference on Information Integration and Web-based Applications & Services (iiWAS '14)*. ACM, New York, NY, USA, 224–228. <https://doi.org/10.1145/2684200.2684314>
- [39] Vivek Nair, Tim Menzies, Norbert Siegmund, and Sven Apel. 2017. Using bad learners to find good configurations. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017, Paderborn, Germany, September 4-8, 2017*. 257–267. <https://doi.org/10.1145/3106237.3106238>
- [40] Vivek Nair, Tim Menzies, Norbert Siegmund, and Sven Apel. 2018. Faster discovery of faster system configurations with spectral learning. *Automated Software Engineering* (2018), 1–31.
- [41] Vivek Nair, Zhe Yu, Tim Menzies, Norbert Siegmund, and Sven Apel. 2018. Finding faster configurations using flash. *IEEE Transactions on Software Engineering* (2018).
- [42] Jeho Oh, Don S. Batory, Margaret Myers, and Norbert Siegmund. 2017. Finding near-optimal configurations in product lines by random sampling. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017, Paderborn, Germany, September 4-8, 2017*. 61–71.
- [43] Jeho Oh, Paul Gazzillo, and Don S. Batory. 2019. *t*-wise coverage by uniform sampling. In *Proceedings of the 23rd International Systems and Software Product Line Conference, SPLC 2019, Volume A, Paris, France*. 15:1–15:4.
- [44] Terence Parr, Kerem Turgutlu, Christopher Csiszar, and Jeremy Howard. 2018. Beware Default Random Forest Importances. last access: july 2019.
- [45] Juliana Alves Pereira, Hugo Martin, Mathieu Acher, Jean-Marc Jézéquel, Goetz Botterweck, and Anthony Ventresque. 2019. Learning Software Configuration Spaces: A Systematic Literature Review. [arXiv:arXiv:1906.03018](https://arxiv.org/abs/1906.03018)
- [46] Quentin Plazar, Mathieu Acher, Gilles Perrouin, Xavier Devroey, and Maxime Cordy. 2019. Uniform Sampling of SAT Solutions for Configurable Systems: Are We There Yet?. In *ICST 2019 - 12th International Conference on Software Testing, Verification, and Validation*. Xian, China, 1–12. <https://hal.inria.fr/hal-01991857>
- [47] Adam Porter, Cemal Yilmaz, Atif M Memon, Douglas C Schmidt, and Bala Nataraajan. 2007. Skoll: A process and infrastructure for distributed continuous quality assurance. *IEEE Transactions on Software Engineering* 33, 8 (2007), 510–525.
- [48] Rodrigo Queiroz, Thorsten Berger, and Krzysztof Czarnecki. 2016. Towards predicting feature defects in software product lines. In *Proceedings of the 7th International Workshop on Feature-Oriented Software Development*. ACM, 58–62.
- [49] Faiza Samreen, Yehia Elkhatib, Matthew Rowe, and Gordon S Blair. 2016. Daleel: Simplifying cloud instance selection using machine learning. In *NOMS 2016-2016 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 557–563.
- [50] Atri Sarkar, Jianmei Guo, Norbert Siegmund, Sven Apel, and Krzysztof Czarnecki. 2015. Cost-Efficient Sampling for Performance Prediction of Configurable Systems (T). *IEEE*, 342–352.
- [51] Norbert Siegmund, Alexander Grebhahn, Sven Apel, and Christian Kastner. 2015. Performance-influence Models for Highly Configurable Systems. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2015)*. 284–294.
- [52] Norbert Siegmund, Sergiy S. Kolesnikov, Christian Kästner, Sven Apel, Don S. Batory, Marko Rosenmüller, and Gunter Saake. 2012. Predicting performance via automated feature-interaction detection. In *ICSE*. 167–177.
- [53] Norbert Siegmund, Marko Rosenmüller, Martin Kuhlemann, Christian Kästner, and Gunter Saake. 2008. Measuring non-functional properties in software product line for product derivation. In *2008 15th Asia-Pacific Software Engineering Conference*. IEEE, 187–194.
- [54] George W Snedecor and Wititiam G Cochran. 1989. *Statistical methods*, 8thEdn. Ames: Iowa State Univ. Press Iowa (1989).
- [55] Charles Song, Adam Porter, and Jeffrey S Foster. 2013. iTree: efficiently discovering high-coverage configurations using interaction trees. *IEEE Transactions on Software Engineering* 40, 3 (2013), 251–265.
- [56] Klaas-Jan Stol and Brian Fitzgerald. 2018. The ABC of Software Engineering Research. *ACM Trans. Softw. Eng. Methodol.* 27, 3, Article 11 (Sept. 2018), 51 pages. <https://doi.org/10.1145/3241743>
- [57] Paul Temple, Mathieu Acher, Battista Biggio, Jean-Marc Jézéquel, and Fabio Roli. 2018. Towards Adversarial Configurations for Software Product Lines. *arXiv preprint arXiv:1805.12021* (2018).
- [58] Paul Temple, Mathieu Acher, Jean-Marc Jézéquel, and Olivier Barais. 2017. Learning Contextual-Variability Models. *IEEE Software* 34, 6 (2017), 64–70. <https://doi.org/10.1109/MS.2017.4121211>
- [59] Paul Temple, José Angel Galindo Duarte, Mathieu Acher, and Jean-Marc Jézéquel. 2016. Using Machine Learning to Infer Constraints for Product Lines. In *Software Product Line Conference (SPLC)*. Beijing, China. <https://hal.inria.fr/hal-01323446>
- [60] Chris Thornton, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2013. Auto-WEKA: Combined selection and hyperparameter optimization of classification algorithms. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 847–855.
- [61] Thomas Thüm, Sven Apel, Christian Kästner, Ina Schaefer, and Gunter Saake. 2014. A Classification and Survey of Analysis Strategies for Software Product Lines. *Comput. Surveys* (2014).
- [62] Pavel Valov, Jianmei Guo, and Krzysztof Czarnecki. 2015. Empirical comparison of regression methods for variability-aware performance prediction. In *SPLC'15*.
- [63] Pavel Valov, Jean-Christophe Petkovich, Jianmei Guo, Sebastian Fischmeister, and Krzysztof Czarnecki. 2017. Transferring performance prediction models across different hardware platforms. In *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering*. ACM, 39–50.
- [64] András Vargha and Harold D Delaney. 2000. A critique and improvement of the CL common language effect size statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics* 25, 2 (2000), 101–132.
- [65] Mahsa Varshosaz, Mustafa Al-Hajjaji, Thomas Thüm, Tobias Runge, Mohammad Reza Mousavi, and Ina Schaefer. 2018. A classification of product sampling for software product lines. In *Proceedings of the 22nd International Conference on Systems and Software Product Line - Volume 1, SPLC 2018, Gothenburg, Sweden, September 10-14, 2018*. 1–13. <https://doi.org/10.1145/3233027.3233035>
- [66] Markus Weckesser, Roland Kluge, Martin Pfanemüller, Michael Matthé, Andy Schürr, and Christian Becker. 2018. Optimal reconfiguration of dynamic software product lines based on performance-influence models. In *Proceedings of the 22nd International Conference on Systems and Software Product Line-Volume 1*. ACM, 98–109.
- [67] Dennis Westermann, Jens Happe, Rouven Krebs, and Roozbeh Farahbod. 2012. Automated inference of goal-oriented performance prediction functions. *ACM*, 190–199.
- [68] Cemal Yilmaz, Myra B Cohen, and Adam A Porter. 2006. Covering arrays for efficient fault characterization in complex configuration spaces. *IEEE Transactions on Software Engineering* 32, 1 (2006), 20–34.
- [69] Yi Zhang, Jianmei Guo, Eric Blais, Krzysztof Czarnecki, and Huiqun Yu. 2016. A mathematical model of performance-relevant feature interactions. In *Proceedings of the 20th International Systems and Software Product Line Conference*. ACM, 25–34.
- [70] Wei Zheng, Ricardo Bianchini, and Thu D Nguyen. 2007. Automatic configuration of internet services. *ACM SIGOPS Operating Systems Review* 41, 3 (2007), 219–229.