



HAL
open science

Learning Software Configuration Spaces: A Systematic Literature Review

Juliana Alves Pereira, Hugo Martin, Mathieu Acher, Jean-Marc Jézéquel,
Goetz Botterweck, Anthony Ventresque

► **To cite this version:**

Juliana Alves Pereira, Hugo Martin, Mathieu Acher, Jean-Marc Jézéquel, Goetz Botterweck, et al..
Learning Software Configuration Spaces: A Systematic Literature Review. [Research Report] 1-44,
Univ Rennes, Inria, CNRS, IRISA. 2019. hal-02148791v1

HAL Id: hal-02148791

<https://inria.hal.science/hal-02148791v1>

Submitted on 7 Jun 2019 (v1), last revised 22 Sep 2021 (v2)

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.

Learning Software Configuration Spaces: A Systematic Literature Review

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

HUGO MARTIN, Univ Rennes, Inria, CNRS, IRISA, France

MATHIEU ACHER, Univ Rennes, Inria, CNRS, IRISA, France

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

GOETZ BOTTERWECK, University of Limerick, Lero–The Irish Software Research Centre, Ireland

ANTHONY VENTRESQUE, University College Dublin, Ireland

Most modern software systems (operating systems like Linux or Android, Web browsers like Firefox or Chrome, video encoders like ffmpeg, x264 or VLC, mobile and cloud applications, etc.) are highly-configurable. Hundreds of configuration options, features, or plugins can be combined, each potentially with distinct functionality and effects on execution time, security, energy consumption, etc. Due to the combinatorial explosion and the cost of executing software, it is quickly impossible to exhaustively explore the whole configuration space. Hence, numerous works have investigated the idea of learning it from a small sample of configurations' measurements. The pattern "sampling, measuring, learning" has emerged in the literature, with several practical interests for both software developers and end-users of configurable systems. In this survey, we report on the different application objectives (e.g., performance prediction, configuration optimization, constraint mining), use-cases, targeted software systems and application domains. We review the various strategies employed to gather a representative and cost-effective sample. We describe automated software techniques used to measure functional and non-functional properties of configurations. We classify machine learning algorithms and how they relate to the pursued application. Finally, we also describe how researchers evaluate the quality of the learning process. The findings from this systematic review show that the potential application objective is important; there are a vast number of case studies reported in the literature from the basis of several domains and software systems. Yet, the huge variant space of configurable systems is still challenging and calls to further investigate the synergies between artificial intelligence and software engineering.

Additional Key Words and Phrases: Systematic Literature Review, Software Product Lines, Machine Learning, Configurable Systems

1 INTRODUCTION

End-users, system administrators, software engineers, and scientists have at their disposal thousands of options (a.k.a. features or parameters) to configure various kinds of software systems in order to fit their functional and non-functional needs (execution time, output quality, security, energy consumption, etc). It is now ubiquitous that software comes in many variants and is highly configurable through conditional compilations, command-line options, runtime parameters, configuration files, or plugins. Software product lines (SPLs), software generators, dynamic, self-adaptive systems, variability-intensive systems are well studied in the literature and enter in this class of configurable software systems [8, 12, 13, 32, 58, 71, 81, 95, 98].

From an abstract point of view, a software configuration is simply a combination of options' values. Though customization is highly desirable, it introduces an enormous complexity due to the combinatorial explosion of possible variants. For example, the Linux kernel has 15,000+ options

Authors' addresses: Juliana Alves Pereira, Univ Rennes, Inria, CNRS, IRISA, Rennes, France, juliana.alves-pereira@irisa.fr; Hugo Martin, Univ Rennes, Inria, CNRS, IRISA, Rennes, France, hugo.martin@irisa.fr; Mathieu Acher, Univ Rennes, Inria, CNRS, IRISA, Rennes, France, mathieu.acher@irisa.fr; Jean-Marc Jézéquel, Univ Rennes, Inria, CNRS, IRISA, Rennes, France, Jean-Marc.Jezequel@irisa.fr; Goetz Botterweck, University of Limerick, Lero–The Irish Software Research Centre, Limerick, Ireland, goetz.botterweck@lero.ie; Anthony Ventresque, University College Dublin, Dublin, Ireland, anthony.ventresque@ucd.ie.




Fig. 1. Features, configurations, sample, measurements, and learning.

and most of them can have 3 values: "yes", "no", or "module". Without considering the presence of constraints to avoid some combinations of options, there may be $3^{15,000}$ possible variants of Linux – the estimated number of atoms in the universe is 10^{80} and is already reached with 300 Boolean options. Though Linux is an extreme case, many software systems or projects exhibit a very large configuration space; it has several consequences.

On the one hand, developers struggle to maintain, understand, and test configuration spaces since they can hardly analyze or execute all variants in every possible settings. According to several studies [30, 81], the flexibility brought by variability is expensive as configuration failures represent one of the most common types of software failures. On the other hand, end-users fear software variability and stick to default configurations [115, 120] that may be sub-optimal (e.g., the software system will run very slowly) or simply inadequate (e.g., the quality of the output will be terrible).

Since it is hardly possible to fully explore all software configurations, the use of machine learning techniques is a quite natural and appealing approach. The basic idea is to learn out of a *sample* of configurations' observations and hopefully generalize to the whole configuration space. There are several applications ranging from performance prediction, configuration optimization, software understanding to constraint mining – we will give a more exhaustive list in this survey. For instance, end-users of x264 (a configurable video encoder) can estimate in advance the execution time of the command-line at the center of Fig. 1, since a machine learning model has been crafted to predict the performance of configurations. End-users may want to use the fastest configuration, or know all configurations that meet an objective (e.g., encoding time should be less than 10 seconds). Developers of x264 can be interested in understanding the effects of some options and how options interact.

For all these use-cases, a pattern has emerged in the scientific literature: "*sampling, measuring, learning*". The basic principle is that a procedure is able to learn out of a sample of configurations' measurements (see Fig. 1). Specifically, many software configuration problems can actually be framed as statistical machine learning problems under the condition a sample of configurations' observations is available. For example, the prediction of the performance of individual configurations can be formulated as a regression problem; appropriate learning algorithms (e.g., CART) can then be used to predict performance of untested, new configurations. In this respect, it is worth noticing the dual use of *feature* in the software or machine learning fields: features either refer to

software features (*a.k.a.* configuration options) or to variables a regressor aims to relate. A way to reconcile and visualize both is to consider a configuration matrix as depicted in Fig. 1. Each row is a configuration together with observations. In the example of Fig. 1, the first configuration has `no_cabac` set to `False` value and `ref` set to `9` value while the encoding time is 3.1876 seconds. We can use a sample of configurations to train a machine learning model (a regressor) with predictive variables being command-line parameters of `x264`. Unmeasured configurations could then be predicted.

Learning software configuration spaces is, however, not a pure machine learning problem and there are a number of specific challenges to address at the intersection of software engineering and artificial intelligence. For instance, 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 can have an important computational cost and is hard to engineer especially when measuring non-functional aspects of software; (3) meanwhile, we expect that the sample is representative to the whole population of valid configurations otherwise the learning algorithm may hardly generalize to the whole configuration space. The general problem is to find the right strategy to decrease the cost of labelling software configurations while minimizing the prediction errors. From an empirical perspective, one can also wonder to what extent learning approaches are effective for real-world software systems present in numerous domains.

While several studies have covered different aspects of configurable systems over the last years, there has been no secondary study (such as systematic literature reviews) that identifies and catalogs individual contributions for machine learning configuration spaces. Thus, there is no clear consensus on what techniques are used to support the process, including which quantitative and qualitative properties are considered and how they can be measured and evaluated, as well as how to select a significant sample of configurations and what is an ideal sample size. This stresses the need for a secondary study to build knowledge from combining findings from different approaches and present a complete overview of the progress made in this field. To achieve this aim, we conduct a *Systematic Literature Review* (SLR) [46] to identify, analyze and interpret all available important research in this domain. We systematically review research papers in which the process of sampling, measuring, and learning configuration spaces occurs – more details about our research methodology are given in Section 2. Specifically, we aim of synthesizing evidence to answer the following six research questions:

- *RQ1. What are the concrete applications of learning software configuration spaces?*
- *RQ2. Which sampling methods are adopted when learning software configuration spaces?*
- *RQ3. Which techniques are used to gather measurements of functional and non-functional properties of configurations?*
- *RQ4. Which learning techniques are used?*
- *RQ5. How are learning-based techniques validated?*
- *RQ6. What are the limitations faced by the current techniques and open challenges that need attention in the future?*

To address *RQ1*, we analyze the application objective of the study (*i.e.*, why they apply learning-based techniques). It would allow us to assess whether the proposed approaches are applicable. With respect to *RQ2*, we investigate which sampling methods are used in the literature. With respect to *RQ3*, we give an in-depth view of how each study measure a sample of configurations. Next, *RQ4* systematically identifies learning techniques used in the literature for exploring the SPL configuration space. In addition, *RQ5* follows identifying which sampling design and evaluation metrics are used for evaluation. Finally, analyzing existing techniques allows identifying evidence

about their maturity and limitations, such as which domains are missing or have not been considered, addressed by *RQ6*.

By answering these questions, we make the following five contributions:

- (1) We identified six main different application areas: *pure prediction, interpretability, optimization, dynamic configuration, evolution, and mining constraints*.
- (2) We provide a framework classification of four main stages used for learning: *Sampling, Measuring, Learning, and Validation*.
- (3) We describe 23 high-level sampling methods, 5 measurement strategies, 51 learning techniques, and 50 evaluation metrics used in the literature. As case study, we identify 71 real-world configurable systems targeting several domains, and functional and non-functional properties. We relate and discuss the learning and validation techniques with regards to their application objective.
- (4) We identify a set of open challenges faced by the current approaches, in order to guide researchers and practitioners to use and build appropriate solutions.
- (5) We build a Web repository to make our SLR results publicly available for the purpose of reproducibility and extension.

Overall, the findings of this SLR reveal that there is a significant body of work specialized in learning software configurable systems with an important application in terms of software technologies, application domains, or goals. There is a wide variety in the considered sampling or learning algorithms as well as in the evaluation process, mainly due to the considered subject systems and application objectives. Practitioners and researchers can benefit from the findings reported in this SLR as a reference when they select a learning technique for their own settings. To this end, this review provides a classification and catalog of specialized techniques in this field.

The rest of the paper is structured as follows. In Section 2, we describe the research protocol used to conduct the SLR. In Section 3, we categorize a sequence of key learning stages used by the ML state-of-the-art literature to explore highly configurable systems. In Section 4, we discuss the research questions. In Section 5, we discuss the threats to the validity of our SLR. In Section 6, we describe similar secondary studies and indicate how our survey differs from them. Finally, in Section 7, we present the conclusions of our work.

2 THE REVIEW METHODOLOGY

We followed the SLR guidelines by Kitchenham and Charters [46] to systematically investigate the use of learning techniques for exploring the SPL configuration space. In this section, we present the SLR methodology that covers three main phases: *planning the review, conducting the review* and *reporting the review*. We report the details about each phase so that readers can assess their rigor and completeness, and reproduce our findings.

2.1 Planning the Review

The steps involved in *planning the review* are: identification of the need for a review, specification of the research questions, and development of a review protocol.

The need for a systematic review. The main goal of this SLR is to systematically investigate and summarize the state-of-the-art of the research concerning learning techniques in the context of software configurable systems. The purpose of conducting this SLR has partially been addressed in the introduction and was motivated by the lack of a systematic study carried on this topic. According to [46] the findings of an SLR is expected to provide a valuable overview of the status of the field to the community through the summarization of existing empirical evidence supported by current scientific studies. The outcomes of such an overview can identify whether, or under what

conditions, the proposed learning approaches can support various use-cases around configurable systems and be practically adopted (e.g., for which context a specific learning technique is much suitable). By mean of this outcome, we can detect the limitations in current approaches to properly suggest areas for further investigation.

The research questions. The goal of this SLR is to answer the following main research question: *What studies have been reported in the literature on learning software configuration spaces since the introduction of Software Product Lines in the early 1990s [43] to date (2019)?* However, this question is too broad, so we derived the six sub-questions defined in Section 1, so as to focus on specific aspects. *RQ1* classifies the papers with regards to their application objective, i.e., for which particular task the approach is suited and useful. We can group studies into similar categories and then compare them. It is also of interest to identify the practical motivations behind learning approaches. We verified whether the authors indicated a specific application for their approach; otherwise, we classified the approach as pure prediction. *RQ2–RQ5* seek to understand key steps of the learning process. *RQ2* reviews the set of sampling methods used in the literature. *RQ3* describes which subject software systems, application domains, and functional and non-functional properties of configurations are measured and how the measurement process is conducted. *RQ4* classifies the set of learning-based techniques used in the literature. *RQ5* aims to characterize the evaluation process used by researcher, including the sample design and supported evaluation metric(s). Finally, addressing these questions will allow us to answer *RQ6*. *RQ6* identifies trends and challenges in the current state-of-the-art approaches, as well as analysis their maturity to summarize our findings and propose future works.

The review protocol. We searched for all relevant papers published up to May 31st 2019. The search process involved the use of 5 scientific digital libraries¹: IEEE Xplore Digital Library², ACM Digital Library³, Science Direct⁴, Springer-link⁵, and Scopus⁶. These search engines were selected because they are known as the top five preferred on-line databases in the software engineering literature [35]. We restricted the search to publication titles and abstracts. However, the library Springer-link only enables a full-text search. Therefore, we first used the full-text option to generate an initial set of papers (the results were stored in a .bib file). Then, we created a script to perform an expert search in the title and abstract over these results.

Each author of this paper had specific roles when performing this SLR. Pereira applied the search string to the scientific databases and exported the results (i.e., detailed information about the candidate papers) into a spreadsheet. After this stage, papers were selected based on careful reading of the titles and abstracts (and if necessary checking the introduction and conclusion). Each identified candidate paper in accordance with the selection criteria defined in Section 2.2 were identified as potentially relevant. When Pereira decided that a paper was not relevant, she provided a short rationale why the paper should not be included in the study. In addition, another researcher checked each excluded paper at this stage. To minimize potential biases introduced into the selection process, any disagreement between researchers were put up for discussion between all authors until a consensus agreement was obtained. This step was done in order to check that all relevant papers were selected.

¹We decided not to use Google Scholar due to search engine limitations, such as the very strict size of the search string.

²<http://ieeexplore.org>

³<http://dl.acm.org>

⁴<http://www.sciencedirect.com>

⁵<http://link.springer.com>

⁶<http://www.scopus.com>

| Term | Keywords |
|------------------------|---|
| Product Line | product line, configurable (system, software), software configurations, configuration of a software, feature (selection, configuration) |
| Learning Techniques | learning techniques, (machine, model, statistical) learning |
| Performance Prediction | performance (prediction, model, goal), (software, program, system) performance, (prediction of, measure) non-functional properties |
| Predict | predict, measure, transfer learning, optimal (configuration, variant), adaptation rules, constraints, context |
| Medicine | gene, medicine, disease, patient, biology, diagnosis, molecular, health, brain, biomedical |

Table 1. Keywords used to build the search strings.

The search in such databases is known to be challenging due to different search limitations, e.g. different ways of constructing the search string. Thus, apart from the automatic search, we also consider the use of snowballing [111] as a complementary approach. Through snowballing, we searched for additional relevant primary studies by following the references from all preliminary selected studies (plus excluded secondary studies). As we published some works related to the topic of the survey, we used our knowledge and expertise to complement the pool of relevant papers.

During the data extraction stage, each paper was assigned to one researcher. Pereira coordinated the allocation of researchers to tasks based on the availability of each researcher. The researcher responsible for extracting the data of a specific selected paper applied the snowballing technique to the correspondent paper. Pereira applied the snowballing technique for excluded secondary studies. Each primary study was then assigned to another researcher for review. Once all the primary studies were reviewed, the extracted data was compared. Whenever there were any discrepancies either about the data reported or about the list of additional selected papers derived from the snowballing process, we again resolved the problem through discussions among all authors.

2.2 Conducting the Review

The steps involved in *conducting the review* are: definition of the search string, specification of the selection criteria, and specification of the data extraction process.

The search string. According to Kitchenham et al. [46] there is no silver bullet for identifying good search strings since, very often, the terminology used in the field is not standardized. When using broad search terms, a large number of irrelevant papers may be found in the search which makes the screening process very challenging. The search string used in this review was first initiated by selecting an initial list of relevant publications by using our expertise in the field.

We identified in the title and abstract the major terms to be used for systematic search of the primary studies. Then, we searched for synonyms related to each major term. Next, we performed several test searches with alternative links between keywords through the different digital libraries. The results from the test searches were continuously discussed among the authors to refine the search string until we were fully satisfied with the capability of the string to detect as much of the initial set of relevant publications as possible. Following this iterative strategy and after a series of test executions and reviews, we obtained Table 1 that structures the set of search terms and keywords.

Specifically, Table 1 shows the *term* we are looking for and related synonyms that we considered as *keywords* in our search. Keywords associated to Product Line allow us to include studies that

focus on configurable systems. By combining keywords associated to Learning Techniques and Performance Prediction, we can find studies that focus on the use of learning-based techniques for exploring the variability space. In addition, keywords associated to Predict (most specific term) allow us to focus on the application objective of such works. We decided to include the keywords associated to Predict so as to identify the context of the study and have a more restricted number of primary studies. Otherwise, the keywords (Product Line AND (Learning Techniques OR Performance Prediction)) return a broad number of studies, *e.g.* studies addressing the use of learning techniques for product line testing or configuration guidance. In addition, we used keywords from Medicine to exclude studies in this field from our search. The final result is the following search string:

(Product Line AND (Learning Techniques OR
Performance Prediction) AND Predict) AND NOT Medicine

The terms Product Line, Learning Techniques, Performance Prediction, and Predict are represented as a disjunction of the keywords in Table 1. The search string format we used was slightly modified to meet the search engine requirements of the different scientific databases. For example, the scientific library Science Direct limits the size of the search string. Thus, when searching in this library, we had to split the search string to generate an initial set of papers. Then, we created a script to perform an expert search of all keywords in the title and abstract over these results, *i.e.* we made every effort to ensure that the search strings used were logically and semantically equivalent to the original string in Table 1. The detailed search strings used in each digital search engine are provided in the Web supplementary material [1].

The selection criteria. The selection of studies was conducted by applying a set of selection criteria for retrieving a relevant subset of publications. First, we only selected papers published up to May 31st 2019 that satisfied all of the following three *Inclusion Criteria* (IC):

- IC1 The paper is available on-line and in English;
- IC2 The paper should be about *configurable software systems*.
- IC3 The paper deals with techniques to statistically learn data from a sample of configurations (see Section 4.1). When different extensions of a paper were observed, *e.g.*, an algorithm is improved by parameter tuning, we intentionally classified and evaluated them as separate primary studies for a more rigorous analysis.

Moreover, we excluded papers that satisfied at least one of the following four *Exclusion Criteria* (EC):

- EC1 Introductions to special issues, workshops, tutorials, conferences, conference tracks, panels, poster sessions, as well as editorials and books.
- EC2 Short papers (less than or equal to 4 pages) and work-in-progress.
- EC3 Pure artificial intelligence papers.
- EC4 Secondary studies, such as literature reviews, articles presenting lessons learned, position or philosophical papers, with no technical contribution. However, the references of these studies were read in order to identify other relevant primary studies for inclusion through the snowballing technique (see Section 2.1). Moreover, we consider secondary studies in the related work section.

We find it useful to give some examples of approaches that were *not* included:

- the use of sampling techniques without learning (*e.g.*, the main application is testing or model-checking a software product line). That is, the sample is not used to train a machine learning model but rather for reducing the cost of verifying a family of products. For a review

on sampling for product line testing, we refer to [19, 50, 54, 56, 102]. We also discuss the complementary between the two lines of work in Section 4.6;

- the use of state-of-the-art recommendations and visualization techniques for configuration guidance (e.g., [69] and [59]) and optimization methods based on evolutionary algorithms (e.g., [28] and [82]) since a sample of configurations' measurements is not considered.
- the use of learning techniques to predict the existence of a software defect or vulnerability based on source code analysis (e.g., in [74] and [94], features do not refer to configurable systems, instead features refer to properties or metrics of the source code).

The data extraction. The data extraction process was conducted using a structured extraction form in Google Sheets⁷ to synthesize all data required for further analyze in such a way that the research questions can be answered. In addition, Google Sheets allow future contributions to be online updated by shareholders. First, all candidate papers were analyzed regarding the selection criteria. The following data were extracted from each retrieved study:

- Date of search, scientific database, and search string.
- Database, authors, title, venue, publication type (i.e., journal, conference, symposium, workshop, or report), publisher, pages, and publication year.
- Inclusion criteria IC1, IC2, and IC3 (yes or no)?
- Exclusion criteria EC1, EC2, and EC3 (yes or no)?
- Selected (yes or no)? If not selected, justification regarding exclusion.

Once the list of primary studies was decided, each selected publication was then read very carefully and the content data for each selected paper was captured and extracted in a second form. The data extraction aimed to summarize the data from the selected primary studies for further analysis of the research questions and for increasing confidence regarding their relevance. All available documentation from studies served as data sources, such as thesis, websites, tool support, as well as the communication with authors (e.g., emails exchanged). The following data were extracted from each selected paper:

- *RQ1:* Scope of the approach. We classified the approach according to the following six categories: *pure prediction, interpretability of configurable systems, optimization, dynamic configuration, mining constraints, and evolution.*
- *RQ2:* Sampling technique(s).
- *RQ3:* Information about subject systems (i.e., reference, name, domain, number of features, and valid configurations) and the measurement procedure. We collected data about the measured (non-)functional properties and the adopted strategies of measurement.
- *RQ4:* Short summary about the adopted learning techniques.
- *RQ5:* Evaluation metrics and sample designs used by approaches for the purpose of training and validating machine learning models.
- *RQ6:* Description of the main challenges and open issues raised by the authors of each selected study. We captured challenges and open issues from the future work, conclusion, or threat to validity sections of a primary study. In the end, we grouped gaps and open challenges according to the learning phase and the scope they address.

2.3 Reporting the Review

The paper selection process is shown in Figure 2. For identifying the candidate primary studies, we followed the review protocol described in Section 2.1. As a first step, we defined our SLR scope. Then, we used our expertise in the field to obtain an initial pool of relevant studies. Based on this

⁷<https://www.google.com/sheets/about/>




Fig. 2. Flow of the paper selection process: papers retrieved from each digital library and details of the selection phases.

effort, we defined the search string (see Section 2.2). As a second step, we applied the search string to the scientific digital libraries. At the end of step 2, the initial search from all sources resulted in a total of 1,627 candidate papers, which includes the 21 papers from our initial pool of relevant studies. Figure 2 shows the number of papers obtained from each digital library.

As a third step, after removing all duplicated papers (285), we carried out the selection process at the content level. During this step, 1,301 papers were excluded, yielding a total of 41 selected papers for inclusion in the review process. A fourth step of the filtering was performed to select additional relevant papers through the snowballing process. This step considered all included papers, as well as removed secondary studies, which resulted in the inclusion of 28 additional papers. This resulted to the selection of 69 primary papers for data extraction. The Web supplementary material [1] provides the results of the search procedure from each of these steps.

3 SURVEY PATTERN: SAMPLING, MEASURING, LEARNING

Understanding how the system behavior varies across a large number of variants of a configurable system is essential for supporting end-users to choose the desirable product. It is also useful of developers in charge of maintaining such software systems. In this context, machine learning-based techniques have been widely considered to predict configurations' behavior and assist stakeholders in making informed decisions. Throughout our surveying effort, we have observed that such approaches follow a 4-stage process: (1) sampling; (2) measuring; (3) learning; and (4) validation. The stages are sketched in Fig. 3. The dashed-line boxes denote the inputs and outputs of each stage. The process starts by building and measuring an initial sample of valid configurations. The set of valid configurations in an SPL is predefined at design time through variability models usually expressed as a feature model [71]. Then, these measurements are used to learn a prediction model. Prediction models help stakeholders to better understand characteristics of




Fig. 3. Employed ML stages to explore SPL configuration spaces.

complex configurable software systems. They try to describe the behavior of all valid configurations. Finally, the validation step computes the accuracy of the prediction model. In addition, some works use active learning [26, 29, 63, 66, 110, 113, 121] to improve the sample in each interaction based on previous accuracy results until it reaches a configuration that has a satisfactory accuracy. Next, we describe in detail each step.

Sampling. Decision makers may decide to select or deselect features to customize a system. Each feature can have an effect on the system non-functional properties. The quantification of the non-functional properties of each individual feature is not enough in most cases, as unknown feature interactions among configuration options may cause unpredictable measurements. Interactions occur when combinations among features share a common component or require additional component(s). Thus, understanding the correlation between feature selections and system non-functional properties is important for stakeholders to be able to find an appropriated system variant that meets their requirements. In Fig. 3, let $C = \{C_1, C_2, \dots, C_n\}$ be the set of n valid configurations, and $C_i = \{f_1, f_2, \dots, f_m\}$ with $f_j \in \{0, 1\}$ a combination of m selected (*i.e.*, 1) and deselected (*i.e.*, 0) features. A straightforward way to determine whether a specific variant meets the requirements is to measure its target non-functional property P and repeat the process for all C variants of a system, and then *e.g.* search for the cheapest configuration C_i with $C_i \in C$. However, it is usually unfeasible to benchmark all possible variants, due to the exponentially growing configuration space. ML techniques address this issue making use of a small measured *sample* $S_C = \{s_1, \dots, s_k\}$ of configurations, where $S_C \subseteq C$, and the number of samples k and the prediction error ϵ are minimal. With the promise to balance measurement effort and prediction accuracy, several sample strategies have been reported in the literature (see Table 3 in Section 4.2). For example, Siegmund et al. [84–87] explore several ways of sampling configurations, in order to find the most accurate prediction model. Moreover, several authors [26, 29, 63, 66, 110, 113, 121] have tried to improve the prediction power of the model by updating an initial sample based on information gained from the previous set of samples through active learning (see Section 4.4). The sample might be partitioned into training, testing and validation sets which are used to train and validate the prediction model (see Section 4.5).

Measuring. This stage measures the set of non-functional properties $\{p_1, \dots, p_l\}$ of a configuration sample $S_C = \{s_1, \dots, s_k\}$, where $p_1 = \{p_1(s_1), \dots, p_1(s_k)\}$. Non-functional properties are measured either by *execution, simulation, static analysis, user feedback* or *synthetic measurements*.

Execution consists of executing the configuration samples and monitoring the measurements of non-functional properties at runtime. Although execution is much more precise, it may incur in unacceptable measurement costs since it is often not possible to create suddenly potentially important scenarios in the real environment. To overcome this issue, some approaches have adopted measurement by simulation. Instead of measuring out of real executions of a system which may result in high costs or risks of failure, *simulation* learns the model using offline environmental conditions that approximate the behavior of the real system faster and cheaper. The use of simulators allows stakeholders to understand the system behavior during early development stages and identify alternative solutions in critical cases. Moreover, simulators can be programmed offline which eliminates any downtime in online environments. In addition to execution and simulation, *static analysis* infers measurements only by examining the code, model, or documentation. For example, the non-functional property cost can be measured as the required effort to add a feature to a system under construction by analyzing the system cycle evolution, such as the number of lines of code, the development time, or other functional size metrics. Although static analysis may not be always accurate, it is much faster than collecting data dynamically by execution and simulation. Moreover, partial configurations can be also measured. Finally, instead of measuring configurations statically or dynamically, some authors also make use of either *user feedback* (UF) or *synthetic measurements* (SM). In contrast to static and dynamic measurements, both approaches do not rely on systems artifacts. *User feedback* relies only on domain experts knowledge to label configurations (e.g., whether the configuration is acceptable or not). *Synthetic measurements* are based on the use of learning techniques to generate artificial (non-)functional values to configurable systems [90]. Researchers can use the THOR generator to mimic and experiment with properties of real-world configurable systems (e.g., performance distributions, feature interactions).

Learning. In this stage, we learn a prediction model based on a given sample of measured configurations $P(S_C)$ to infer the behavior of non-measured configurations $P(C - S_C)$. The sampling set S_C is divided into a *training set* S_T and a *validation set* S_V , where $S_C = S_T + S_V$. The training set is used as input to learn a prediction model, *i.e.* describe how configuration options and their interactions influence the behavior of a system. For parameter tuning, interactive sampling, and active learning, the training set is also partitioned into training and testing sets (see Section 4.4).

Some authors [39–41, 105, 106] applied transfer learning techniques to accelerate the learning process. Instead of building the prediction model from scratch, transfer learning reuses the knowledge gathered from samples of other relevant related sources to a target source. It uses a regression model that automatically captures the correlation between target and related systems. Correlation means the common knowledge that is shared implicitly between the systems. This correlation is an indicator that there is a potential to learn across the systems. If the correlation is strong, the transfer learning method can lead to a much more accurate and reliable prediction model more quickly by reusing measurements from other sources.

Validating. The validation stage quantitatively analysis the quality of the sample S_T for prediction using an evaluation metric on the validation set S_V . To be practically useful, an ideal sample S_T should result in a (i) low prediction error; (ii) small model size; (iii) reasonable measurement effort. The aim is to find as few samples as possible to yield an understandable and accurate ML model in short computation time. In Section 4.5, we detail the evaluation metrics used by the selected primary studies to compute accuracy.

Overall, exploring the configuration space based on a small sample of configurations is a critical step since in practice, the sample may not contain important feature interactions nor reflect the system real behavior accurately. To overcome this issue, numerous learning approaches have been

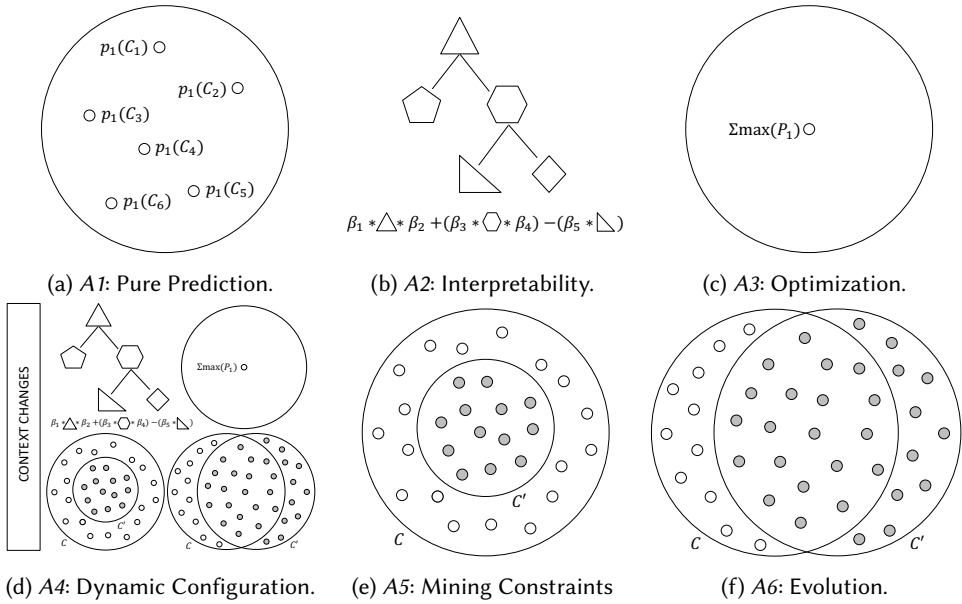


Fig. 4. Study application objective.

proposed in the last years. Next, we analyze the existing literature by investigating the more fine-grained characteristics of these four learning stages.

4 RESULTS AND DISCUSSION OF THE RESEARCH QUESTIONS

In this section, we discuss the answers to our research questions defined in Section 1. In Section 4.1, we identify the main goal of the learning process. Next, in Sections 4.2, 4.3, 4.4 and 4.5 we deeply analyze how each study address each learning stage defined in Section 3 to accomplish the goals described in Section 4.1. Finally, Section 4.6 presents a set of limitations and open challenges in order to point out possible future directions in this domain.

4.1 RQ1: What are the concrete applications of learning software configuration spaces?

In this section, we analyze the application objective of the selected studies since learning may have different practical interests and motivations. Learning techniques have been used in the literature to target six different scenarios (see Figure 4 and Table 2). We describe each individual application and give prominent examples of papers entering in this category.

A1 Pure Prediction. The aim is to accurately predict labels of unmeasured configurations. Labels can be qualitative (e.g., whether the software configuration has a defect) or quantitative (e.g., the execution time in seconds of a software configuration). The outcome is to associate through prediction some properties to all configurations of the space (see Figure 4a). Guo et al. [27] is a seminal paper with respect to the use of statistical learning for predicting performances. In this scenario, other factors such as the model comprehension and the computation time are less important. In some engineering contexts, the sole prediction of a property of a configuration has limited practical interest *per se* and is sometimes used as a basis for targeting other applications (e.g., configuration optimization).

- A2 *Interpretability of configurable systems.* Understanding the correlation between configuration options and system quality is important for a wide variety of tasks, such as optimization, program comprehension and debugging. To this end, these studies aim at learning an accurate model that is fast to compute and simple to interpret (*i.e.*, easy for developers to get an understandable overview of how configuration options interact, see Figure 4b). For example, Kolesnikov et al. [48] explore how so-called performance-influence models quantify options' influences and can be used to explain the performance behavior of a configurable system as a whole.
- A3 *Optimization.* Instead of labeling all configurations, optimization approaches aim at finding a (near-)optimal valid configuration to best satisfy requirements (see Figure 4c). According to Ochoa et al. [64], there are three types of stakeholders' requirements: resource constraints (threshold such as response time < 1 hour), stakeholders' preferences (*e.g.*, security is extremely more preferable and relevant than response time) and optimization objectives (*e.g.*, minimization of response time). Although the specification of requirements may reduce the configuration space [2, 9, 21, 65, 76], searching for the most appropriate configuration is still an overwhelming task due to the combinatorial explosion. Learning approaches exposed in *e.g.* [61, 63, 66] propose a recursive search method to interactively add new samples to train the prediction model until it reaches a sample with an acceptable accuracy.
- A4 *Dynamic Configuration.* There are many dynamic systems (*e.g.*, robotic systems) that aim to manage run-time adaptations of software to react to (uncertain) environmental changes. Without self-adaptation, requirements would be violated. There are several works that explicitly define a set of adaptation rules in the variability model during design-time. However, anticipating all contextual changes and defining appropriate adaptation rules earlier is often hard to domain engineers due to the huge variability space and the uncertainty of how the context may change at run-time. Therefore, approaches classified in this group use learning techniques to constantly monitor the environment to detect contextual changes that require the system to adapt (see Figure 4d). It learns the influences of contexts online in a feedback loop under time and resource constraints to reach one of the previous defined application objective. For example, learning techniques are used in the dynamic scenario to support the synthesis of contextual-variability models, including logical constraints [49, 98]. Other approaches [37, 41, 73, 79, 109] use learning techniques with the goal of finding a configuration that is optimal and consistent with the current, dynamic context.
- A5 *Mining Constraints.* In a configurable system, not all combinations of options' values are possible (*e.g.*, some options are mutually exclusive). Variability models are used to precisely define the space of valid configurations, typically through the specification of logical constraints among options. However, the identification of constraints is a difficult task and it is easy to forget a constraint leading to configurations that do not compile, crash at run-time, or do not meet a particular resource constraint or optimization goal. To overcome this issue, learning techniques can be used to discover additional constraints that would exclude unacceptable configurations. These studies seek an accurate and complete set of constraints to restrict the space of possible configurations. Finally, it creates a new variability model by adding the identified constraints to the original model (see Figure 4e). Therefore, the aim is to accurately remove invalid configurations that were never derived and tested before. Mining constraints approaches work with qualitative properties, such as *video quality* [97, 100] and *defects* [7, 24, 49, 116].
- A6 *Evolution.* In the evolution scenario, a configurable system will inevitably need to adapt to satisfy real-world changes in external conditions, such as changes in requirements, design and performance improvements, and changes in the source code. Thus, new configuration

| App. | Reference |
|------|--|
| A1 | [15, 17, 27, 29, 40, 42, 47, 52, 75, 80, 84–87, 90, 91, 93, 104, 117–119] |
| A2 | [20, 23, 39, 48, 90, 92, 105] |
| A3 | [6, 10, 18, 22, 25, 26, 36, 38, 55, 60–63, 66, 67, 78, 88–90, 96, 101, 106, 110, 113, 114, 120, 121] |
| A4 | [14, 20, 37, 41, 49, 73, 79, 83, 90, 98, 109] |
| A5 | [3, 7, 24, 49, 77, 90, 97, 98, 100, 116] |
| A6 | [83, 90, 120] |

Table 2. Applicability of selected primary studies. A1: Pure Prediction; A2: Interpretability of Configurable Systems; A3: Optimization; A4: Dynamic Configuration; A5: Mining Constraints; A6: SPL Evolution.

options become available and valid, while existing configurations may become obsolete and invalid (see Figure 4f). Consequently, the new variability model structure may influence certain non-functional properties. In this scenario, it is important to make sure that the learning stage is informed by *evolution* about changes and the set of sample configurations is readjusted accordingly with the new variability model by excluding invalid configurations and considering the parts of the configuration space not yet covered. In this case, the measurements for each configuration that includes a new feature or a new interaction is updated.

The idea of "sampling, measuring, learning" has a wide application objective and can be used for supporting developers or end-users in six main different tasks: *Pure Prediction, Interpretability of Configurable Systems, Optimization, Dynamic Configuration, Mining Constraints, and SPL Evolution*. It is also possible to combine different tasks (e.g., mining constraints for supporting dynamic configuration [49, 98]).

4.2 RQ2: Which sampling methods are adopted when learning software configuration spaces?

In this section, we analyze the set of sampling methods used by learning-based techniques in the literature. Table 3 shows the sample methods adopted by each study. The first column is about the method and the second column identifies the study reference(s). There are 23 high-level sample methods documented in the literature. Next, we describe the particularities of the most used sample methods.

Random sampling. Several studies have used random sampling [3, 27, 29, 37, 39, 41, 42, 61–63, 66, 84, 97, 98, 100, 104, 105, 109, 118] with different notions of randomness.

Guo et al. [27] consider four sizes of random samples for training: N , $2N$, $3N$, and M , where N is the number of features of a system, and M is the number of minimal valid configurations covering each pair of features. They choose size N , $2N$, and $3N$, because measuring a sample whose size is linear in the number of features is likely feasible and reasonable in practice, given the high cost of measurements by execution (see Section 4.3). Valov et al. and Guo et al. [29, 104, 105] use a random sample to train, but also to cross-validate their machine learning model. Several works [61, 63, 66, 118] seek to determine the number of samples in an adaptive, progressive sampling manner and a random strategy is usually employed. Nair et al. [61, 63] and Jehoo et al. [66] aim at optimizing a configuration. At each iteration, they randomly add an arbitrary number of configurations to learn a prediction model until they reach a model that exhibits a desired satisfactory *accuracy* (see Section 4.4). They consider several sizes of samples from tens to thousands of configurations. To focus on a reduced part of the configuration space, Nair et

| Sample Method | Reference |
|--------------------------------------|---|
| Random | [3, 6, 14, 26, 27, 29, 37, 39, 41, 42, 61–63, 66, 67, 84, 89, 90, 96–98, 100, 101, 104, 105, 109, 118, 119] |
| Knowledge-wise heuristic | [86, 87, 89, 90, 114] |
| Feature-coverage heuristic | [26, 42, 52, 80, 84–87, 116, 117] |
| Feature-frequency heuristic | [23, 26, 52, 80, 84–88] |
| Family-based simulation | [91] |
| Multi-start local search | [36] |
| Plackett-Burman design | [26, 84] |
| central composite design | [26] |
| D-optimal design | [26] |
| Breakdown | [110] |
| Sequence type trees | [49] |
| East-west sampling | [62] |
| Exemplar sampling | [62] |
| Constrained-driven sampling | [24] |
| Diameter uncertainty strategy | [121] |
| Historical dataset of configurations | [78, 106] |
| Latin hypercube sampling | [10, 38, 113] |
| Neighborhood sampling | [40, 73] |
| Input-based clustering | [18] |
| Distance-based sampling | [25, 42] |
| Genetic sampling | [38, 55, 77] |
| Interaction tree discovery | [93] |
| Arbitrarily chosen | [7, 14, 15, 20, 60, 75, 79, 92] |

Table 3. Sample methods reported in the literature.

al. [63] and Jehoo et al. [66] determine statistically significant parts of the configuration space that contribute to good performance through active learning. In order to have a more representative sample, Valov et al. [105] adopted stratified random sampling. This sampling strategy exhaustively divides a sampled population into mutually exclusive subsets of observations before performing actual sampling.

Prior works [27, 29, 105] relied on the random selection of features to create a configuration, followed by a filter to eliminate invalid configurations (*a.k.a.* pseudo-random sampling). Walker’s alias sampling [105] is an example of pseudo-random sampling. Quasi-random sampling (*e.g.*, sobol sampling) is similar to pseudo-random sampling, however they are specifically designed to cover a sampled population more uniformly [6, 105]. However, pseudo-random sampling may result in too many invalid configurations, which makes this strategy inefficient. To overcome this issue, several works [3, 37, 39, 41, 63, 66, 97, 98, 100, 109, 118] use solver-based sampling techniques (*a.k.a.*, true random sampling).

Sampling and heuristics. Instead of randomly choosing configurations as part of the sample, several heuristics have been developed. The general motivation is to better cover features and features’ interactions as part of the sample. The hope is to *better capture the essence of the configuration space with a lower sampling size*. We describe some heuristics hereafter.

Knowledge-wise heuristic. This heuristic selects a sample of configurations based on its influence on the target non-functional properties. Siegmund et al. [86, 87] sampling method measures each

feature in the feature model plus all known feature interactions defined by a domain expert. Experts detect feature interactions by analyzing the specification of features, implementation assets, and source code, which require substantial domain knowledge and exhaustive analysis. SPLCoqueror⁸ provides to stakeholders an environment in which they can document and incorporate known feature interactions. For each defined feature interaction, a single configuration is added to the set of sample for measurement. THOR [90] is a generator for synthesizing synthetic yet realistic variability models where users (researchers) can specify the number of interactions and the degree of interactions.

Feature-coverage heuristic. To automatically detect all first order feature interactions, Siegmund et al. [84–87] use a pair-wise measurement heuristic. This heuristic assumes the existence of a feature interaction between each pair of features in an SPL. It includes a minimal valid configuration for each pair of features being selected. Pair-wise requires a number of measurements that is quadratic in the number of optional features. Some authors [42, 52, 80] also use a 3-wise feature coverage heuristic to discover interactions among 3 features. Siegmund et al. [85] propose a *3rd-order coverage heuristic* that considers each minimal valid configuration where three features interact pair-wise among them (adopted by [52]). They also propose the idea that there are hot-spot features that represent a performance-critical functionality within a system. These hot-spot features are identified by counting the number of interactions per features from the feature-coverage and higher-order interaction heuristics. Yilmaz et al. [117] adopted even a 4-wise feature coverage heuristic, and Yilmaz et al. [116] a 5 and 6-wise heuristic. As there are n-th order feature coverage heuristics, the sample set might be likely unnecessarily large which increases measurement effort substantially. However, not every generated sample contains features that interacts with each other. Thus, the main problem of this strategy is that it requires prior knowledge to select a proper coverage criterion. To overcome this issue, state-of-the-art approaches might use the interaction-wise heuristic to fix the size of the initial sample to the number of features or potential feature interactions of a system [90].

Feature-frequency heuristic. The feature-frequency heuristic considers a set of valid configurations in which each feature is selected and deselected, at least, once. Sarkar et al. [80] heuristic counts the number of times a feature has been selected and deselected. Sampling stops when the counts of features selected and deselected is, at least, at a predefined threshold. Nair et al. [61] analysis the number of samples required by using the previous heuristic [80] against a rank-based random heuristic. Siegmund et al. [84–88] quantify the influence of an individual feature by computing the delta of two minimal configurations with and without the feature. They then relate to each feature a minimum valid configuration that contains the current feature, which requires the measurement of a configuration per feature. Hence, each feature can exploit the previously defined configuration to compute its delta over a performance value of interest. In addition, to maximize the number of possible interactions, Siegmund et al. [84] also relate to each feature a maximal valid configuration that contains the current feature.

There are several others sampling heuristics, such as *Plackett-Burman design* [26, 84] for reasoning with numerical options; *Breakdown* [110] (random breakdown, adaptive random breakdown, adaptive equidistant breakdown) for breaking down (in different sectors) the parameter space; *Constrained-driven sampling* [24] (constrained CIT, CIT of constraint validity, constraints violating CIT, combinatorial union, unconstrained CIT) to verify the validity of combinatorial interaction testing (CIT) models; and many others (see Table 3).

⁸<http://fosd.de/SPLConqueror>.

Sampling and transfer learning. Jamshidi et al. [39–41] aim at applying transfer learning techniques to learn a prediction model (see Section 4.4). Jamshidi et al. [41] consider a combination of random samples from *target* and *source* systems for training: {0%, 10%, ..., 100%} from the total number of valid configurations of a source system, and {1%, 2%, ..., 10%} from the total number of valid configurations of a target system. In a similar scenario, Jamshidi et al. [39] randomly select an arbitrary number of valid configurations from a system before and after environmental changes (e.g., using different hardware, different workloads, and different versions of the system). In another scenario, Jamshidi et al. [40] use transfer learning to sample. Their sampling strategy, called L2S, exploits common similarities between source and target systems. L2S progressively learns the interesting regions of the target configuration space, based on transferable knowledge from the source.

Arbitrarily chosen sampling. Chen et al. [15] and Murwantara et al. [60] have arbitrarily chosen a set of configurations as their sample is based on their current available resources. Sincero et al. [92] use a subset of valid configurations from a preliminary set of (de)selected features. This sample is arbitrarily chosen by domain experts based on the use of features which will probably have a high influence on the properties of interest. In the context of investigating temporal variations, Samreen et al. [79] consider on-demand instances at different times of the day over a period of seven days with a delay of ten minutes between each pair of runs. In a similar context, Duarte et al. [20] and Chen et al. [14] also sample configurations under different workloads (e.g., active servers and requests per second) at different times of the day. An important insight is that there are engineering contexts in which the sampling strategy is imposed and can hardly be controlled.

All configurations (no sampling). Sampling is not applicable for four of the selected primary studies [17, 47, 48, 120], mainly for experimental reasons. For example, Kolesnikov et al. [47, 48] consider all valid configurations in their experiments and use an established learning technique to study and analyze the trade-offs among prediction error, model size, and computation time of performance-prediction models. For the purpose of their study, they were specifically interested to explore the evolution of the model properties to see the maximum possible extent of the corresponding trade-offs after each iteration of the learning algorithm. So, they performed a whole-population exploration of the largest possible learning set (i.e., all valid configurations). In a similar scenario, Kolesnikov et al. [47] explored the use of control-flow feature interactions to identify potentially interacting features based on detected interactions from performance prediction techniques using performance measurements). Therefore, they also performed a whole-population exploration of all valid configurations.

Reasoning about configurations validity. Sampling is realized either out of an enumerated set of configurations (e.g., the whole ground truth) or a variability model (e.g., a feature model). The former usually assumes that configurations of the set are logically valid. The latter is more challenging, since picking a configuration boils down to resolve a satisfiability or constraint problem.

Acher et al. [3] and Temple et al. [97, 98, 100] encoded variability models as *Constraint Programming* (CSP) by using the Choco solver, while Weckesser et al. [109] and Siegmund et al. [90] employed SAT solvers. Constraint solver may produce clustered configurations with similar features due to the way solvers enumerate solutions (i.e., often the sample set consists of the closest k valid configurations). Therefore, these strategies do not guarantee true randomness as in pseudo-random sampling. Moreover, using CSP and SAT solvers to enumerate all valid configurations are often impractical [57, 72]. Thus, Jehoo et al. [66] encoded variability models as *Binary Decision Diagrams* (BDDs) [4], for which counting the number of valid configurations is straightforward. Given the number of valid configurations n , they randomly and uniformly select the k^{th} configuration, where

$k \in \{1 \dots n\}$, *a.k.a.* randomized true-random sampling. Kaltenecker et al. [42] perform a comparison among pseudo-random sampling, true-random sampling, and randomized true-random sampling.

Reasoning with numerical options. Ghamizi et al. [25] transform numeric and enumerated attributes into alternative Boolean features to be handle as binary option. Temple et al. [97, 98, 100] adopted random sampling of numeric options, *i.e.* real and integer values. First, their approach randomly selects a value for each feature within the boundaries of its domain. Then, it propagates the values to other features with a solver to avoid invalid configurations. In a similar scenario, Siegmund et al. [84] and Grebhahn et al. [26] adopted pseudo-random sampling of numeric options. Siegmund et al. [84] claim that it is very unusual that numeric options have value ranges with undefined or invalid holes and that constraints among numeric options appear rarely in configurable systems. Grebhahn et al. [26] adopted different reasoning techniques over binary and numeric options, then they compute the Cartesian product of the two sets to create single configurations used as input for learning. In the scenario of reasoning with numerical options, Amand et al. [7] arbitrarily select equidistant parameters values.

Though random sampling is a widely used baseline, numerous other sampling algorithms and heuristics have been devised and described. There are different trade-offs to find when sampling configurations (1) minimization of invalid configurations due to constraints' violations among options; (2) minimization of the cost (*e.g.*, size) of the sample; (3) generalization of the sample to the whole configuration space. The question of a one-size-fits-all sampling strategy remains open and several factors are to be considered (targeted application, subject systems, functional and non-functional properties, presence of domain knowledge, etc.).

4.3 RQ3: Which techniques are used to gather measurements of functional and non-functional properties of configurations?

The measuring step takes as input a sample of configurations and measures, for each configuration, their functional properties or *non-functional properties* (NFPs). In this section, we investigate how the measurement procedures are technically realized.

Most proposals consider NFPs, such as elapsed time in seconds. In essence, NFPs consist of a *name*, a *domain*, a *value*, and a *unit* [11]. The domain type of a NFP can be either quantitative (*e.g.*, real and integer) or qualitative (*e.g.*, string and Boolean). *Quantitative (QT)* properties are typically represented as a numeric value, thus they can be measured on a metric scale, *e.g.*, the configuration is executed in 13.63 seconds. *Qualitative (QL)* properties are represented using an ordinal scale, such as low (−) and high (+); *e.g.*, the configuration produces a high video quality.

As described in Section 3, measurements can be obtained through five different strategies: *execution* (EX), *static analysis* (SA), *simulation* (SI), *user feedback* (UF), and *synthetic measurements* (SM). Fig. 5 shows that automated execution is by far the most used technique to measure configuration properties.

In Table 4, we present the subject systems used in the literature together with NFPs. The first column identifies the references. The second and third columns describe the name and domain of the system, respectively. The fourth column points out the measured NFP(s). There are 71 real-world SPLs documented in the literature. They are of different sizes, complexities, implemented in different programming languages (C, C++, and Java), varying implementation techniques (conditional compilation and feature-oriented programming), and from several different application domains (*e.g.*, operating systems, video encoder, database system) and developers (both academic and industrial). Therefore, they cover a broad spectrum of scenarios. It is important to mention that the

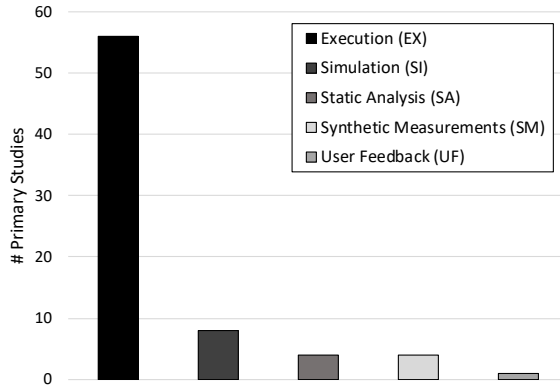


Fig. 5. Strategies applied to measure the sample of configurations.

same subject systems may differ in the number of features, feature interactions, and in the number of valid configurations – the experimental setup is simply different.

Table 4. Subject Systems supported in the literature.

| References | Name | Domain | Non-Functional Properties |
|---|----------------|----------------------|---|
| [7] | Thingiverse’s | 3D printer | defects |
| [113] | IBM WebSphere | Application server | throughput |
| [29, 48, 84] | Clasp | ASP solver | response time |
| [121] | SNW | Asset management | area and throughput |
| [18] | Binpacking | Binpacking algorithm | execution time and accuracy |
| [40] | XGBoost | Boosting algorithms | training time |
| [83] | SaaS system | Cloud computing | response time |
| [18] | Clustering | Clustering algorithm | execution time and accuracy |
| [29, 48, 84, 91] | AJStats | Code analyzer | response time |
| [39, 40, 63, 84] | SaC | Code analyzer | I/O time, response time |
| [42] | POLLY | Code optimizer | runtime |
| [24] | Libssh | Combinatorial model | defects |
| [24] | Telecom | Communication system | defects |
| [27, 29, 42, 48, 62, 63, 66, 80, 84, 85, 87, 98, 118, 119, 121] | LLVM | Compiler | memory footprint, performance, response time, code complexity, compilation time |
| [52] | Compressor SPL | Compression library | compression time, memory usage and compression ratio |
| [42] | 7Z | Compression library | Compression time |
| [29, 42, 48, 61, 84] | LRZIP | Compression library | compressed size, compression time, compilation time |
| [87] | RAR | Compression library | code complexity |
| [105] | XZ | Compression library | compression time |
| [86–88, 91] | ZipMe | Compression library | memory footprint, performance, code complexity, time |
| [60] | WordPress | Content management | CPU power consumption |
| [86–89] | LinkedList | Data structures | memory footprint, performance, maintainability, binary size |
| [87] | Curl | Data transfer | code complexity |

Table 4. Subject Systems supported in the literature.

| References | Name | Domain | Non-Functional Properties |
|---|------------------|------------------------|---|
| [61, 87] | Wget | Data transfer | memory footprint, code complexity |
| [106] | Actian Vector | Database system | runtime |
| [41] | Apache Cassandra | Database system | latency |
| [27, 29, 42, 48, 61, 62, 66, 80, 84–89, 98, 118, 119] | Berkeley DB | Database system | I/O time, memory footprint, performance, response time, code complexity, maintainability, binary size |
| [89] | FAME-DBMS | Database system | maintainability, binary size, performance |
| [93, 106, 117, 120] | MySQL | Database system | defects, throughput, latency |
| [106] | Postgres | Database system | throughput, latency |
| [86–88] | Prevayler | Database system | memory footprint, performance |
| [27, 29, 39, 47, 61, 62, 80, 85–87, 87, 88, 98, 105] | SQLite | Database system | memory footprint, performance, response time, code complexity, runtime |
| [15] | StockOnline | Database system | response time |
| [10] | Kafka | Distributed systems | throughput |
| [25] | DNN | DNNs algorithms | accuracy of predictions |
| [3] | Curriculum vitae | Document | number of pages |
| [3] | Paper | Document | number of pages |
| [20] | RUBIS | E-commerce application | response time |
| [91] | EMAIL | E-mail client | time |
| [47] | MBED TLS | Encryption library | response time |
| [110] | SAP ERP | Enterprise Application | response time |
| [63] | noc-CM-log | FPGA | CPU power consumption, runtime |
| [63] | sort-256 | FPGA | area, throughput |
| [23] | E-Health System | Health | response time |
| [29, 42, 84, 98] | HIPAcc | Image processing | response time |
| [17] | Disparity SPL | Image processing | energy consumption |
| [86–88] | PKJab | Instant messenger | memory footprint, performance |
| [36] | IBM ILOG CPLEX | Integer solver | runtime |
| [110] | SPECjbb2005 | Java Server | response time, throughput |
| [101] | WEKA | Learning algorithm | accuracy of predictions |
| [18] | SVD | Linear algebra | execution time and accuracy |
| [63] | Trimesh | Mesh solver | iterations, response time |
| [91] | MBENCH | Micro benchmark | time |
| [73, 116] | ACE+TAO system | Middleware software | defects |
| [86–88] | SensorNetwork | Network simulator | memory footprint, performance |
| [109] | Simonstrator | Network simulator | latency |
| [121] | NoC | Network-based system | energy and runtime |
| [18] | Helmholtz 3D | Numerical analysis | execution time and accuracy |
| [18] | Poisson 2D | Numerical analysis | execution time and accuracy |
| [86–88, 92] | Linux kernel | Operating system | memory footprint, performance |
| [40] | DNN | Optimization algorithm | response time |
| [55] | Art system | Paint | user likeability |
| [26] | Multigrid system | Equations solving | time of each interaction |
| [41] | CoBot System | Robotic system | CPU usage |
| [42, 84, 98] | JavaGC | Runtime environment | response time |
| [96] | Robot | Runtime environment | energy consumption and execution time |

Table 4. Subject Systems supported in the literature.

| References | Name | Domain | Non-Functional Properties |
|--|------------------|------------------------|--|
| [22, 114] | Hand | SAT solver | runtime |
| [22, 114] | Indu | SAT solver | runtime |
| [22, 114] | Rand | SAT solver | runtime |
| [36] | SAPS | SAT solver | runtime |
| [36, 39] | SPEAR | SAT solver | runtime, response time |
| [78] | QWS dataset | Services | availability, throughput, success-ability, reliability, compliance, best practice, documentation |
| [75] | BusyBox | Software suite | defect, process metrics |
| [49] | Plant automation | Software-intensive SoS | defects |
| [18] | Sort | Sort algorithm | execution time |
| [42, 48, 61, 84, 98] | DUNE | Stencil code | response time |
| [48, 61, 84, 98] | HSMGP | Stencil code | response time, runtime |
| [38, 40, 41, 61, 63, 118, 119] | Apache | Stream processing | latency, throughput, performance |
| [24] | Concurrency | Testing problem | defects |
| [79] | VARD on EC2 | Text manager | response time |
| [24] | Aircraft | Toy example | defects |
| [91] | ELEVATOR | Toy example | time |
| [24] | WashingMachine | Toy example | defects |
| [86–88] | Violet | UML editor | memory footprint, performance |
| [97, 98, 100] | MOTIV | Video encoder | video quality |
| [42] | VP9 | Video encoder | encoding time |
| [27, 29, 39, 42, 48, 61–63, 66, 80, 84, 85, 87, 98, 105, 118, 119] | x264 | Video encoder | CPU power consumption, encoding time, Peak Signal to Noise Ratio, response time, code complexity, video quality, performance |
| [98] | OpenCV | Video tracking | performance |
| [77] | C60 and MX300 | Virtual environment | defects |
| [6] | Amazon EC2 | Web cloud service | performance |
| [27, 29, 48, 61, 62, 66, 80, 84, 85, 98, 117, 120] | Apache | Web server | response rate, response time, workload, defects, throughput |
| [67] | Stock Brokerage | Web system | throughput, response time |
| [93] | vsftpd | FTP daemon | defects |
| [93] | ngIRCd | IRC daemon | defects |

Next, we detail the particularities of NFPs. Specifically, we describe how the measurement is performed, what process and strategies are adopted to avoid biases in the results, and also discuss the cost of measuring.

Time. The time spent by a software configuration to realize a task is an important concern and has been intensively considered under different flavors and terminologies (see Fig. 6). Siegmund et al.[84] invested more than two months (24/7) for measuring *response time* of all configurations of different subject systems (Dune MGS, HIPAcc, HSMGP, JavaGC, SaC, x264). For each system they use a different configuration of hardware. The configurations’ measurements are reused in many papers, mainly for evaluating the proposed learning process [29, 62, 99]. Chen et al. [15] used

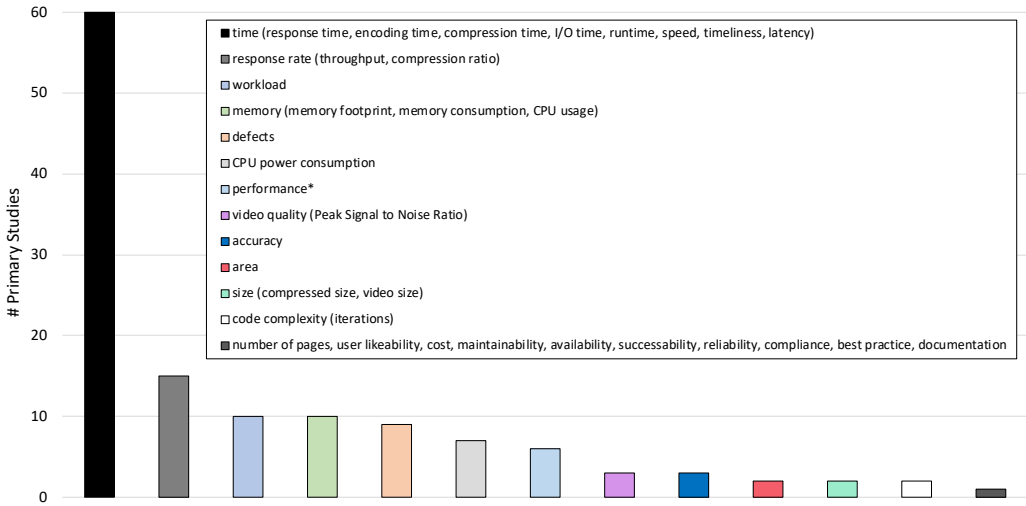


Fig. 6. Non-functional properties measured in the literature.

*Each application domain or even each system has special definition of performance. These works did not specify the NFP(s) that refer to performance.

a simulator to measure the response time of a sequence of 10 service requests by a client to the server. They use two implementation technologies, CORBA and EJB.

Time is a general notion and can be refined according to the application domain. For instance, Kolesnikov et al. [48] considered *compression time* of a file on Lrzip (Long Range ZIP), a file archiver written. The measurements were conducted on a dedicated server and repeated multiple times to control measurements noise. Kaltenecker et al. [42] considered 7-ZIP, a file archiver written in C++, and measured the compression time of the Canterbury corpus⁹.

In another engineering context, several works [27, 42, 48, 61, 62, 104, 118] measured the *encoding time* of an input video over 1,152 valid configurations of x264. x264 is a configurable system for encoding video streams into the H.264/MPEG-4 AVC format. As benchmark, the Sintel trailer (735 MB) is used and an encoding from AVI is considered. Kaltenecker et al. [42] measured the encoding time of a short piece of the Big Buck Bunny trailer over 216,000 valid configurations of VPXENC (VP9), a video encoder that uses the VP9 format.

Latency has caught several attention. Jamshidi et al. [38] and [41] measure the average latency (*i.e.*, how fast it can respond to a request) of three stream processing applications on Apache Storm (WordCount, RollingSort, SOL) over a window of 8 minutes and 2 hours, respectively. Jamshidi et al. [41] also measure the average latency of the NoSQL database system on Apache Cassandra over a window of 10 minutes. Jamshidi et al. [38] performed the measurements on a multi-node cluster on the EC2 cloud. Jamshidi et al. [41] performed the same measurement procedure reported for CPU usage. The measurements used by Nair et al. [61, 63] were derived from Jamshidi et al. [38]. Weckesser et al. [109] measure the latency of transferred messages over an adaptive Wireless Sensor Networks (WSNs) via simulation. 100 fully-charged nodes were distributed randomly onto a square region for each simulation run. Aken et al. [106] measure the latency for two OLTP DBMSs (MySQL v5.6, Postgres v9.3) over three workloads (The Yahoo Cloud Serving Benchmark (YCSB), TPC-C, and Wikipedia) during five minutes observation periods. The OLTP DBMS are deployed on

⁹<https://corpus.canterbury.ac.nz/>

m4.large instances with 4 vCPUs and 16 GB RAM on Amazon EC2. They also consider the total execution time of the OLAP DBMS (Actian Vector v4.2) over the TPC-H workload on m3.xlarge instances with 4 vCPUs and 15 GB RAM on Amazon EC2. All of the training data was collected using the DBMSs' default isolation level.

Sharifloo et al. [83] measures time throughout the evolution of a configurable system (an SPL). In an SPL evolution scenario, the set of measured configurations may include a removed feature or violate changed constraints. In this case, configurations are removed from the sample. Moreover, a modified feature implies to recompute the configuration measurements.

Other NFPs (CPU power consumption, CPU usage, size, etc.) Beyond time, other quantitative properties have been considered.

Nair et al. [61] measured the compressed size of a specific input over 432 valid configurations of Lrzip, a compression program optimized for large files. Siegmund et al. [86–88] measured the *memory footprint* of nine configurable system: LinkedList, Prevayler, ZipMe, PKJab, SensorNetwork, Violet, Berkeley DB, SQLite, and Linux kernel. Nair et al. [63] and Zuluaga et al. [121] used LLVM, a configurable modular compiler infrastructure. Footprint is measured as the binary size of a compiled configuration.

Several works [17, 37, 60, 63, 98, 121] measured *CPU power consumption*. Murwantara et al. [60] used a simulator to compute CPU power consumption over a web service under different loads. They used a kernel-based virtual machine to conduct the measurements based on several combinations of HTTP servers and variant PHP technologies connected to a MySQL database system. They divided the experiment into blocks of 10 seconds for 100 increasing stages. In the first 10 seconds, they produce loads of one user per second, and in the next 10 seconds, they produce the load of two users per seconds and so on. This results in 1-100 users per period of 10 seconds. Couto et al. [17] propose a static learning approach of CPU power consumption. Their approach assures that the source code of every feature from a configuration is analyzed only once. To prove its accuracy, they measured at runtime the worst-case CPU power consumption to execute a given instruction on 7 valid configurations of the disparity SPL [108]. They repeated the measurements 200 times for each configuration using the same input. Since their goal was to determine the worst case CPU power consumption, they removed the outliers (5 highest and lowest values) and retrieved the highest value from every configuration for validation proposes.

Jamshidi et al. [41] use simulation measurements. Jamshidi et al. [41] have repeatedly executed a specific robot mission to navigate along a corridor offline in a simulator and measured performance in terms of *CPU usage* on the CoBot system. To understand the power of transfer learning techniques, they consider several simple hardware changes (*e.g.*, processor capacity) as well as severe changes (*e.g.*, local desktop computer to virtual machines in the cloud)¹⁰. Each measurement took about 30 seconds. They measured each configuration of each system and environment 3 times.

To overcome the cost of measuring realistic (non-)functional properties, Siegmund et al. [90] proposed a tool, called Thor, to generate artificial and realistic synthetic measurements, based on different distribution patterns of property values for features, interactions, and configurations from a real-world system. Jamshidi et al. [37] adopted Thor to synthetically measure the property energy consumption for a robot to complete a mission consisting of randomly chosen tasks within a map.

Qualitative properties. Instead of measuring a numerical value, several papers assign a qualitative value to configurations. In Acher et al. [3], a variant is considered as acceptable or not thanks to an automated procedure. In Temple et al. [97, 100], a quantitative value is originally measured and then transformed into a qualitative one through the definition of a threshold. In [75], a variant is considered as acceptable or not based on the static evolution historical analysis of commit messages

¹⁰For a complete list of hardware/software variability, see <https://github.com/pooyanjamshidi/transferlearning>

(i.e., they identify a defect by searching for the following keywords: *bug*, *fix*, *error*, and *fail*). The learning process then aims to predict the class of the configuration – whether configurations are acceptable or not, as defined by the threshold.

Software *defects* are considered in [7, 24, 49, 73, 116, 117]. Yilmaz et al. [116] and Porter et al. [73] characterize defects of the ACE+TAO system. Yilmaz et al. [116] test each supposed valid configuration on the Red Hat Linux 2.4.9-3 platform and on Windows XP Professional using 96 developer-supplied regression tests. In [73], developers currently run the tests continuously on more than 100 largely uncoordinated workstations and servers at a dozen sites around the world. The platforms vary in versions of UNIX, Windows, Mac OS, as well as to real-time operating systems. To examine the impact of masking effects on Apache v2.3.11-beta and MySQL v5.1, Yilmaz et al. [117] grouped the test outcomes into three classes: passed, failed, and skipped. They call this approach a *ternary-class fault characterization*. Gargantini et al. [24] interest is in comparing the defect detection capability on different sample heuristics (see Section 4.2), while Amand et al. [7] interest is in comparing the accuracy of several learning algorithms to predict whether a configuration will lead to a defect. In the dynamic configuration scenario, Krismayer et al. [49] use event logs (via simulation) from a real-world automation SoS to mine different types of constraints according to real-time defects.

Finally, Martinez et al. [55] consider a 5-point *user likeability* scale with values ranging from 1 (strong dislike) to 5 (strong like). In this work, humans have reviewed and labeled configurations.

Accuracy of measurements. In general, measuring NFPs (e.g., time) is a difficult process since several confounding factors should be controlled. The need to gather measures over numerous configurations exacerbates the problem.

Weckesser et al. [109] mitigated the construct threat of the inherent randomness and repeated all runs five times with different random seeds. Measurements started after a warm-up time of 5 minutes. Kaltenecker et al. [42] measured each configuration between 5 to 10 times until reaching a standard deviation of less than 10%. Zhang et al. [118] repeated the measurements 10 times.

To investigate the influence of measurement errors on the resulting model, Kolesnikov et al. [48] and Duarte et al. [20] conducted a separate experiment. They injected measurement errors to the original measurements and repeated the learning process with *polluted* datasets. Then, they compared the prediction error of the noisy models to the prediction error of the original models to see the potential influence of measurement errors. For each subject system, Kolesnikov et al. [48] repeated the learning process five times for different increasing measurement errors.

Dynamic properties are susceptible to measurement errors (due non-controlled external influences) which may bias the results of the measuring process. To account for measurement noise and be subject to external influences, these properties need to be measured multiple times on dedicated systems. Thus, the total measurement cost to obtain the whole data used in the experiments is overly expensive and time-consuming (e.g., Kaltenecker et al. [42] spent multiple years of CPU time). According to Lillacka et al. [52], there are a warm-up phase followed by multiple times runs; and the memory must be set up large enough to prevent disturbing effects from the Garbage Collector, as well as all operations, must be executed in memory so that disk or network I/O will also produce no disturbing effects. Most of the works only consider the variability of the subject system, while they use static inputs and hardware/software environments. Therefore, the resulting model may not properly characterize the performance of a different input or environment, since most of the properties (e.g., CPU power consumption and compression time) are dependent of the input task and the used hardware/software. Consequently, hardware/software must also be taken into account as dependent variables as considered by Jamshidi et al. [39, 41].

There are some properties that are much accurate, because e.g. they are not influenced by the used hardware, such as footprint. Siegmund et al. [87] parallelized the measurements of footprint

on three systems and used the same compiler. Moreover, footprint can be measured quickly only once, without measurement bias.

Cost of measurements. The cost of observing and measuring software can be important, especially when multiple configurations should be considered. The cost can be related to computational resources needed (in time and space). It can also be related to human resources involved in labelling some configurations [55].

Zuluaga et al. [121] and Nair et al. [63] measure the quantitative NFP *area* of a field-programmable gate array (FPGA) platform consisting of 206 different hardware implementations of a sorting network for 256 inputs. They report that the measurement of each configuration is very costly and can take up to many hours. Porter et al. [73] characterization of defects for each configuration ranges from 3 hours on quad-CPU machines to 12-18 hours on less powerful machines. Yilmaz et al. [116] took over two machine years to run the total of 18,792 valid configurations. In Siegmund et al. [88], a single measurement of memory footprint took approximately 5 minutes. Temple et al. [100] report that the execution of a configuration to measure video quality took 30 minutes on average. They used a grid computing to distribute the computation and scale the measurement for handling 4,000+ configurations. The average time reported by Murwantara et al. [60] to measure CPU power consumption of all sampled configurations was 1,000 seconds. The authors set up a predefined threshold to speed up the process.

There are fifteen main NFPs supported in the literature. Some of them are less costly, such as *code complexity*, which can be measured statically by analyzing the number of code lines [88]. As a result, the measurements can be parallelized and quickly done only once. Otherwise, dynamic properties, such as *CPU power consumption* and *response time* are directly related to hardware and external influences. While *CPU power consumption* might be measured under different loads over a predefined threshold time, the property *response time* is much costly as a threshold can not be defined and the user does not know how long it will take. To ease the labeling of features, some works use synthetic NFPs and statistical analysis strategies. As the same NFP may be measured in different ways (e.g., video quality can be measured either by *user feedback* or *execution* of a program to automatically attribute labels to features). Practitioners need to find a sweet spot to have accurate measurements with a small sample.

Depending on the subject system and application domain (see the dataset of [84]), there are more favorable cases with only a few seconds per configuration. However, even in this case, the overall cost can quickly become prohibitive when the number of configurations to measure is too high. In [109], all training simulation runs took approximately 412 hours of CPU time.

Numerous qualitative and quantitative properties of configurations are measured mainly through the use of automated software procedures. For a given subject system and its application domain, there may be more than one measure (e.g., CPU power consumption and video quality for x264). Time is the most considered performance measure and is obtained in the literature through either execution, simulation, static analysis, or synthetic measurement. The general problem is to find a good tradeoff between the cost and the accuracy of measuring numerous configurations (e.g., simulation can speed up the measurement process at the price of approximating the real observations).

4.4 RQ4: Which learning techniques are used?

In this section, we aim at understanding which learning techniques were applied and in which context. Table 5 sketches which learning techniques are supported in the literature and what application objective they address. The first column identifies the study reference(s). The second and

third columns identify the name of the learning technique and its application objective, respectively. (Notice that the application objective is related to the scenarios in which each learning technique has been already used in the literature; it means some learning techniques could well be applied for other scenarios in the future).

Supervised learning problems can be grouped into regression and classification problems. In both cases, the goal is to construct machine learning model that can predict the value of the measurement from the features. The difference between the two problems is the fact that the value to predict is numerical for regression and categorical for classification. In the survey, we found a similar dichotomy, depending on the targeted use-case and the NFP of interest.

A *regression problem* is when the output is a real or continuous value, such as time or CPU power consumption. Most of the learning techniques tackle a supervised regression problem.

CART for regression. Several authors [27, 29, 39, 60, 61, 63, 80, 98, 100] use the *Classification And Regression Trees* (CART) technique, to model the correlation between feature selections and performance. The sample is used to build the prediction model. CART recursively partitions the sample into smaller clusters until the performance of the configurations in the clusters are similar. These recursive partitions are represented as a binary decision tree. For each cluster, these approaches use the sample mean of the performance measurements (or even the majority vote) as the local prediction model of the cluster. So, when they need to predict the performance of a new configuration not measured so far, they use the decision tree to find the cluster which is most similar to the new configuration. Each split of the set of configurations is driven by the (de)selection of a feature that would minimize a prediction error.

CART use two parameters to automatically control the recursive partitioning process: *minbucket* and *minsplit*. Minbucket is the minimum sample size for any leaf of the tree structure; and minsplit is the minimum sample size of a cluster before it is considered for partitioning. Guo et al. [27] compute minbucket and minsplit based on the size of the input sample, *i.e.*, if $|S_C| \leq 100$, then $\text{minbucket} = \lfloor \frac{|S_C|}{10} + \frac{1}{2} \rfloor$ and $\text{minsplit} = 2 \times \text{minbucket}$; if $|S_C| > 100$, then $\text{minsplit} = \lfloor \frac{|S_C|}{10} + \frac{1}{2} \rfloor$ and $\text{minbucket} = \lfloor \frac{\text{minsplit}}{2} \rfloor$; the minimum of minbucket is 2; and the minimum of minsplit is 4. It should be noted that CART can also be used for classification problems (see hereafter).

Instead of using a set of empirically-determined rules, Guo et al. [29] combine the previous CART approach [27] with automated resampling and parameter tuning, which they call a data-efficient learning approach (DECART). Using resampling, DECART learns a prediction model by using different sample designs (see Section 6). Using parameter tuning, DECART ensures that the prediction model has been learned using optimal parameter settings of CART based on the currently available sample. They compare three parameter-tuning techniques: random search, grid search, and Bayesian optimization. Westermann et al. [110] also used grid search for tuning CART parameters.

Nair et al. [61, 63, 80] approaches build a prediction model in a progressive way by using CART. They start with a small training sample and subsequently add samples to improve performance predictions based on the model accuracy (see Section 4.5). In each step, while training the prediction model, Nair et al. [61] compare the current accuracy of the model with the previous accuracy from the prior iteration (before adding the new set of configurations to the training set). If the current accuracy (with more data) does not improve the previous accuracy (with lesser data), then the learning reaches a termination criterion (*i.e.*, adding more sample will not result in significant accuracy improvements).

Performance-influence models. Siegmund et al. [84] combines machine learning and sampling heuristics to build so-called performance-influence models. A step-wise linear regression algorithm is used to select relevant features as relevant terms of a linear function and learn their

| Reference | Learning Technique | Applicability |
|---|--|------------------------|
| [117] | Adaptive ELAs, Multi-Class FDA-CIT, Static Error Locating Arrays (ELAs), Ternary-Class FDA-CIT, Test Case-Aware CIT, Traditional CIT | A1 |
| [104] | Bagging | A1 |
| [118] | Fourier Learning of Boolean Functions | A1 |
| [47] | Frequent Item Set Mining | A1 |
| [91] | Graph Family-Based Variant Simulator | A1 |
| [17] | Implicit Path Enumeration Technique (IPET) | A1 |
| [75] | Naive Bayes | A1 |
| [37, 39, 42, 47, 48, 84, 89, 109] | Step-Wise Linear Regression | A1, A2, A3, A4 |
| [3, 27, 29, 39, 49, 60–63, 73, 80, 93, 98, 100, 104, 105, 110, 116, 119, 120] | Classification and Regression Trees (CART) | A1, A2, A3, A4, A5 |
| [90] | Kernel Density Estimation and NSGA-II | A1, A2, A3, A4, A5, A6 |
| [85–88] | Feature’s Influence Delta | A1, A3 |
| [6, 38, 40, 41, 106, 110, 121] | Gaussian Process Model | A1, A3, A4 |
| [15, 52, 60, 79] | Linear Regression | A1, A3, A4 |
| [7, 10, 75, 101, 104] | Random Forest | A1, A3, A5 |
| [7, 22, 97, 104] | Support Vector Machine | A1, A3, A5 |
| [7, 75, 93] | C4.5 (J48) | A1, A5 |
| [23, 92] | Covariance Analysis | A2 |
| [39] | Multinomial Logistic Regression | A2 |
| [20] | K-Plane Algorithm | A2, A4 |
| [78] | AdaRank | A3 |
| [60] | Bagging Ensembles of CART, Bagging Ensembles of MLPs | A3 |
| [55] | Data Mining Interpolation Technique | A3 |
| [106] | Factor Analysis, k-means, Ordinary Least Squares | A3 |
| [96, 110] | Genetic Programming (GP) | A3 |
| [110] | Kriging | A3 |
| [18] | Max-Apriori Classifier, Exhaustive Feature Subsets Classifiers, All Features Classifier, Incremental Feature Examination classifier | A3 |
| [67] | Quick Optimization via Guessing | A3 |
| [36] | Random Online Adaptive Racing (ROAR), Sequential Model-based Algorithm Configuration (SMAC) | A3 |
| [96] | Simulated Annealing | A3 |
| [113] | Smart Hill-Climbing | A3 |
| [66] | Statistical Recursive Searching | A3 |
| [25] | Tensorflow and Keras | A3 |
| [101] | Tree-structured Parzen Estimator (TPE) | A3 |
| [26, 79, 110] | Multivariate Adaptive Regression Splines (MARS), Multivariate Polynomial Regression | A3, A4 |
| [79, 114] | Ridge Regression | A3, A4 |
| [7, 60] | Multilayer Perceptrons (MLPs) | A3, A5 |
| [14] | Actor-Critic Learning | A4 |
| [79] | Lasso | A4 |
| [83] | Reinforcement Learning | A4, A6 |
| [24] | CitLab Model | A5 |
| [97] | Evasion Attack | A5 |
| [7] | Hoeffding Tree, K*, kNN, Logistic Model Tree, Logistic Regression, Naive Bayes, PART Decision List, Random Committee, REP Tree, RIPPER | A5 |
| [77] | Pruning Rule-Based Classification (PART) | A5 |

Table 5. Learning techniques reported in the literature. A1: Pure Prediction; A2: Interpretability of Configurable Systems; A3: Optimization; A4: Dynamic Configuration; A5: Mining Constraints; A6: SPL Evolution.

coefficient to explain the observations. In each iterative step, the algorithm selects the sample configuration with strongest influence regarding prediction accuracy (*i.e.*, yields the model's lowest prediction error) until improvements of model accuracy become marginal or a threshold for expected accuracy is reached (below 19%). The algorithm concludes with a backward learning step, in which every relevant feature is tested for whether its removal would decrease model accuracy. This can happen if initially a single feature is selected because it better explains the measurements, but it becomes obsolete by other features (*e.g.*, because of feature interactions) later in the learning process. Linear regression allows them to learn a formula that can be understood by humans. It also makes it easy to incorporate domain knowledge about an option's influence on the formula. However, the complete learning of a model using this technique required from 1 to 5 hours, depending on the size of the learning set and the size of the models. Kolesnikov et al. [48] investigate how significant are the trade-offs among prediction error, model size, and computation time.

Other learning algorithms for regression. Westermann et al. [110] used Multivariate Adaptive Regression Splines (MARS), Genetic Programming (GP), and Kriging. Zhang et al. [118] used Fourier learning algorithm. In all these works, for a set of sample, it verifies if the resulting accuracy is acceptable for stakeholders (*e.g.*, prediction error rate below 10%). While the accuracy is not satisfactory, the process continues by obtaining an additional sample of measured configurations and iterates again to produce an improved prediction model. Sarkar et al. [80] propose a sampling cost metric as a stopping criterion, where the objective is to ensure the most optimal trade-off between measurement effort and prediction accuracy. Sampling stops when the counts of features selected and deselected is, at least, at a predefined threshold.

Chen et al. [15] propose a linear regression approach to describe the generic performance behavior of application server components running on component-based middleware technologies. The model focuses on two performance factors: workload and degree of concurrency. Sincero et al. [92] employ analysis of covariance for identifying factors with significant effects on the response or interactions among features. They aim at proposing a configuration process where the user is informed about the impact of their feature selection on the NFPs of interest.

Murwantara et al. [60] use a set of five ML techniques (*i.e.*, Linear regression, CART, Multilayer Perceptrons (MLPs), Bagging Ensembles of CART, and Bagging Ensembles of MLPs) to learn how to predict the energy consumption of web service systems. They use the WEKA's implementation of these techniques with its default parameters [31].

Learning to rank. Instead of predicting the raw performance value, it can be of interest to predict the rank of a configuration (typically to identify *optimal* configurations, see RQ1).

Jehoo et al. [66] adopt statistical learning techniques to progressively shrink a configuration space and search for near-optimal configurations. First, the approach samples and measures a set of configurations. For each pair of sampled configurations, this approach identifies features that are common (de)selected. Then, it computes the performance influence of each common decision to find the best regions for future sampling. The performance influence measures the average performance over the samples that have the feature selected against the samples that have the feature deselected, *i.e.*, the sample is partitioned by whether a configuration includes a particular feature or not. In addition, Welch's t-test evaluates whether the performance mean of one sample group is higher than the other group with 95% confidence. The most influential decisions are added to the sample. This process continues recursively until they identify all decisions that are statistically certain to improve program performance. They call this a *Statistical Recursive Searching* technique.

Nair et al. [61] compute accuracy by using the mean rank difference measurement (the predicted rank order is compared to the optimal rank order). They demonstrate that their approach can find optimal configurations of a software system using fewer measurements than the approach

proposed by [80]. One drawback with this approach is that it requires a holdout set, against which the current model (built interactively) is compared. To overcome this issue, instead of making comparisons, Nair et al. [63] consider a predefined stopping criterion (budget associated with the optimization process). While the criterion is not met, the approach finds the configuration with the best accuracy and add the configuration to the training set. The model adds the next most promising configuration to evaluate. Consequently, in each step, the approach discards less satisfactory configurations which have a high probability of being dominated, *a.k.a.* active learning. This process terminates when the predefined stopping condition is reached. Nair et al. [63] demonstrate that their approach is much more effective for multi-objective configuration optimization than state-of-the-art approaches [61, 80, 121].

Martinez et al. [55] propose the use of data mining interpolation techniques (*i.e.*, similarity distance, similarity radius, weighted mean) for ranking configurations through user feedback on a configuration sample. They estimate the user perceptions of each feature by computing the value of the chi-squared statistic with respect to the correlation score given by users on configurations.

Transfer learning. The previous approaches assume a static environment (*e.g.*, hardware, workload) and NFP such that learning has to be repeated once the environment and NFP changes. In this scenario, some approaches adopt transfer learning techniques to reuse (already available) knowledge from other relevant sources to learn a performance for a target system instead of relearn a model from scratch [39–41, 105].

Valov et al. [105] investigate the use of transfer learning across different hardware platforms. They used 25 different hardware platforms to understand the similarity on performance prediction. They created a prediction model using CART. With CART, the resulting prediction models can be easily understood by end users. To transfer knowledge from a related source to a target source, they used a simple linear regression model.

Jamshidi et al. [41] use *Gaussian Process Models* (GPM) to model the correlation between source and target sources using the measure of *similarity*. GPM offers a framework in which predictions can be done using mean estimates with a confidence interval for each estimation. Moreover, GPM computations are based on linear algebra which is cheap to compute. This is especially useful in the domain of dynamic SPL configuration where learning a prediction model at runtime in a feedback loop under time and resource constraints is typically time constrained.

Jamshidi et al. [39] combine many statistical and ML techniques (*i.e.*, *Pearson linear correlation*, *Kullback-Leibler divergence*, *Spearman correlation coefficient*, *paired t-test*, *CART*, *step-wise linear regression*, and *multinomial logistic regression*) to identify when transfer learning can be applied. They use CART for estimating the relative importance of configuration options by examining how the prediction error will change for the trained trees on the source and target. To investigate whether interactions across environments will be preserved, they use step-wise linear regression models (*a.k.a.*, performance-influence models). This model learns all pairwise interactions, then it compares the coefficients of the pairwise interaction terms independently in the source and target environments. To avoid exploration of invalid configurations and reduce measurement effort, they use a multinomial logistic regression model to predict the probability of a configuration being invalid, then they compute the correlation between the probabilities from both environments.

Jamshidi et al. [40] propose a sampling strategy, called L2S, that exploits common similarities across environments from [39]. L2S extracts transferable knowledge from the source to drive the selection of more informative samples in the target environment. Based on identifying interesting regions from the performance model of the source environment, it generates and selects configurations in the target environment iteratively.

Classification problem. Temple et al. [98, 100] and Acher et al. [3] use CART to infer constraints to avoid the derivation of invalid (non-acceptable) configurations. CART considers a path in a tree

as a set of decisions, where each decision corresponds to the value of a single feature. The approach creates new constraints in the variability model by building the negation of the conjunction of a path to reach a faulty leaf. They learn constraints among Boolean and numerical options. As an academic example, Acher et al. [3] introduce Vary \LaTeX to guide researchers to meet paper constraints by using annotated \LaTeX sources. To improve the learning process, Temple et al. [97] specifically target low confidence areas for sampling. The authors apply the idea of using an adversarial learning technique, called evasion attack, after a classifier is trained with a Support Vector Machine (SVM). In addition, Temple et al. [98] support the specialization of configurable systems for a deployment at runtime. In a similar context, Safdar et al. [77] infer constraints over multi-SPLs (*i.e.*, they take into account cross-SPLs rules).

Input sensitivity. Several works consider the influence of the input data on the resulting prediction model [18, 22, 25, 26, 36, 52, 101, 114]. Many authors [18, 22, 25, 26, 36, 101, 114] address input sensitivity in algorithm auto-tuning. They use learning techniques to search for the best algorithmic variants and parameter settings to achieve optimal performance for a given input instance. It is well known that the performance of SAT solver and learning methods strongly depends on making the right algorithmic and parameter choices, therefore SATzilla [114] and Auto-WEKA [101] search for the best SAT solver and learning technique for a given input instance. Lillacka et al. [52] treat the variability caused by the input data as the variability of the SPL. As it is not feasible to model every possible input data, they cluster the data based on its relevant properties (*e.g.*, 10kB and 20MB in an or-group for file inputs of a compression library).

Reinforcement learning. Sharifloo et al. [83] use a reinforcement learning technique to automatically reconfigure dynamic SPLs to deal with context changes. In their approach, learning continuously observes measured configurations and evaluates their ability to meet the contextual requirements. Unmet requirements are addressed by learning new adaptation rules dynamically, or by modifying and improving the set of existing adaptation rules. This stage takes software evolution into account to address new contexts faced at run-time.

Numerous statistical learning algorithms are used in the literature to learn software configuration spaces. The most used are standard machine learning techniques, such as polynomial linear regressions, decision trees, or Gaussian process models. The targeted use-case and the engineering context explain the diversity of solutions: either a supervised classification or regression problem is tackled; the requirements in terms of interpretability and accuracy may differ; there is some innovation in the sampling phase to progressively improve the learning. Still the use of others (more powerful) ML techniques such as deep learning, adversarial learning, and even the idea of learning different models (*e.g.*, one for each application objective) that could co-evolve can be further explored in future works.

4.5 RQ5: How are learning-based techniques validated?

In this section, we aim at understanding how the validation process is conducted in the literature.

There are five design strategies documented in the literature to explore the sample data for learning and validation (see Table 6).

Merge. Training, testing, and validation sets are merged. Kolesnikov et al. [48] studied and analyzed the trade-offs among prediction error, model size, and computation time of performance-prediction models. For the purpose of their study, they were specifically interested to explore the evolution of the model properties to see the maximum possible extent of the corresponding trade-offs after each iteration of the learning algorithm. So, they performed a whole-population exploration of the largest possible learning set (*i.e.*, all valid configurations). Therefore, in their validation process,

| Sample Design | Reference |
|---------------------------|--|
| Merge | [15, 42, 48, 63, 66, 84–88, 92] |
| Hold-Out | [3, 6, 7, 10, 14, 20, 22, 25–27, 29, 38, 40, 41, 49, 52, 61, 62, 75, 78, 80, 93, 96–98, 100, 104, 105, 109, 110, 114, 116, 119, 121] |
| Cross-Validation | [18, 20, 29, 55, 60, 77, 79, 117] |
| Bootstrapping | [29, 38, 61, 101, 106] |
| Dynamic Sector Validation | [110] |

Table 6. Sample designs reported in the literature.

training, testing and validation sets are merged. In a similar scenario, other authors [15, 63, 66] also used the whole sample for both, learning and validation, even they considered a small set of valid configurations as the sample. Some studies justify the use of a merge pool due the need to compare different sampling methods. However, training the algorithm on the validation pool may introduce bias to the results of the accurateness of the approach. To overcome this issue, studies have used other well-established ML design strategies.

In particular, Guo et al. [29] study the trade-offs among hold-out, cross-validation, and bootstrapping. For most of the case studies, 10-fold cross-validation outperformed hold-out and bootstrapping. In terms of the running time, although these three strategies usually takes seconds to run, Guo et al. [29] shows that hold-out is the fastest one, and 10-fold cross-validation tends to be faster than bootstrapping.

Hold-Out (HO). Most of the works [3, 6, 7, 27, 29, 38, 40, 41, 49, 61, 62, 80, 97, 98, 100, 104, 105, 109, 110, 116, 121] used a hold-out design. Hold-out splits an input sample S_C into two disjointed sets S_t and S_v , one for training and the other for validation, i.e. $S_C = S_t \cup S_v$ and $S_t \cap S_v = \emptyset$. To avoid bias in the splitting procedure, some works repeated it multiple times. For example, in Valov et al. [104], the training set for each sample size were selected randomly 10 times. For transfer-learning applications [40, 41, 105], the training set comes from samples of the target and related sources, while the validation set comes from samples only from the target source.

Cross-Validation (CV). Cross-validation splits an input sample S_C into k disjointed subsets of the same size, i.e., $S = S_1 \cup \dots \cup S_k$, where $S_i \cap S_j = \emptyset$ ($i \neq j$); each subset S_i is selected as the validation set S_v , and all the remaining $k - 1$ subsets are selected as the training set S_t . Yilmaz et al. [117] used a 5-fold cross-validation to create multiple models from different subsets of the input data. In a 5-fold cross-validation, the sample S is partitioned into 5 subsets (i.e., $S = S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$) of the same size. Each subset is selected as the validation set and all of the remaining subsets form the training set. Most authors [29, 55, 60, 77, 79] relayed on a 10-fold cross-validation. 10-fold cross-validation follows the same idea of a 5-fold cross-validation, producing $k = 10$ groups of training and validation sets. Guo et al. [29] show that a 10-fold cross-validation design does not work well for very small samples, e.g., it did not work for Apache system when the sample size was 9 configurations.

Bootstrapping (BT). Four studies [29, 38, 61, 106] relayed on the bootstrapping design. Bootstrapping relies on random sampling with replacement. Given an input sample S_C with k configurations, bootstrapping randomly selects a configuration C_b , with $1 \leq b \leq k$ and copies it to the training set S_t . However, it keeps C_b in S_C for the next selection. This process is repeated k times. Given the training sample, bootstrapping uses $S_C \setminus S_t$ as the validation set S_v . Nair et al. [61] used a non-parametric bootstrap test with 95% confidence for evaluating the statistical significance of

their approach. In non-parametric bootstrap, the input sample S_C is drawn from a discrete set, *i.e.*, 95% of the resulting sample S_t should fall within the 95% confidence limits about S_C .

Dynamic Sector Validation. Westermann et al. [110] relayed on two types of Dynamic Sector Validation designs: with local prediction error scope (DSL) and with global prediction error scope (DSG). Dynamic Sector Validation decides if a sample configuration $C_i \in S_C$ is part either of the training or testing set based on the sector’s prediction error of the adopted learning techniques (see Section 4.4). In DSL, all sectors have a prediction error that is less than the predefined threshold, while in DSG the average prediction error of all sectors is less than the predefined threshold.

There are studies where a sample design is *not applicable* [17, 23, 24, 37, 39, 47, 73, 89, 90, 113, 120], as well as studies [36, 67, 83, 91, 118] without further details about the sample design. Given the sampling design, evaluation metrics are used to (learn and) validate the resulting prediction model. Table 7 sketches which evaluation metrics are used in the literature. The first column identifies the study reference(s). The second and third columns identify the name of the evaluation metric and its application objective, respectively. Similar to Table 5, here the application objective is related to the scenarios in which each evaluation metric has been already used in the SPL literature. Therefore, in future researches, some metrics may be explored in other scenarios as well. Notice that some studies [83, 88, 92] did not report any detail about the validation process.

Evaluation metrics. State-of-the-art techniques rely on 50 evaluation metrics from which it is possible to evaluate the accuracy of the resulting models in different application scenarios. There are metrics dedicated to supervised classification problems (*e.g.*, precision, recall, F-measure). In such settings, the goal is to quantify the ratio of correct classifications to the total number of input samples. For instance, Temple et al. [100] used precision and recall to control whether acceptable and non-acceptable configurations are correctly classified according to the ground truth. Others [73, 97] use qualitative analysis to identify features with significant effects on defects, and understand feature interactions and decide whether further investigation of features is justified.

There are also well known metrics for regression problems, such as, *Mean Relative Error* (MRE - Equation 1) and *Mean Absolute Error* (MAE - Equation 2). These metrics aim at estimating the accuracy between the exact measurements and the predicted one.

$$MRE = \frac{1}{|S_v|} \sum_{C \in S_v} \frac{|C_{(p_i)} - C_{(p'_i)}|}{C_{(p_i)}} \quad (1)$$

$$MAE = \frac{|C_{(p'_i)} - C_{(p_i)}|}{C_{(p_i)}} \quad (2)$$

Where S_v is the validation set, and $C_{(p_i)}$ and $C_{(p'_i)}$ indicate the exact and predicted value of p_i for $C \in S_v$, respectively.

Contributions addressing learning-to-rank problems develop specific metrics capable of assessing the ability of a learning method to correctly rank configurations. For example, Nair et al. [61, 63] use the error difference between the ranks of the predicted configurations and the true measured configurations. To get insights about when stop sampling, they also discuss the trade-off between the size of the training set and rank difference.

Interpretability metrics. Some metrics are also used to better *understand* configuration spaces and their distributions, for example, *when* to use a transfer learning technique. In this context, Jamshidi et al. [39] use a set of similarity metrics (Kullback-Leibler, Pearson Linear Correlation, Spearman Correlation Coefficient, and Rank Correlation) to investigate the relatedness of the source and target environments. These metrics compute the correlation between predicted and exact values from

| Reference | Evaluation Metric | Applicability |
|--|--|----------------|
| [104] | Closeness Range, Winner Probability | A1 |
| [93, 119] | Coverage | A1 |
| [47] | Jaccard Similarity | A1 |
| [119] | Performance-Relevant Feature Interactions Detection Accuracy | A1 |
| [117] | t-masked metric | A1 |
| [75] | True Positive (TP) Rate, False Positive (FP) Rate, Receiver Operating Characteristic (ROC) | A1 |
| [15, 17, 20, 26, 27, 29, 37, 40, 42, 48, 62, 84–87, 101, 104, 105, 110, 118] | Mean Relative Error (MRE) | A1, A2, A3, A4 |
| [6, 20, 26, 62, 63, 80, 93, 109] | Sampling Cost | A1, A2, A3, A4 |
| [110, 118] | Highest Error (HE) | A1, A3 |
| [10, 22, 25, 38, 40, 41, 52, 55, 60, 66, 91] | Mean Absolute Error (MAE) | A1, A3, A4 |
| [36, 42, 77] | Mann-Whitney U-test | A1, A3, A5 |
| [42, 79] | F-test | A1, A4 |
| [3, 7, 47, 98, 100] | Precision, Recall | A1, A4, A5 |
| [7, 75, 77, 116, 117] | Precision, Recall, F-measure | A1, A5 |
| [23] | GAP | A2 |
| [39] | Kullback-Leibler (KL), Pearson Linear Correlation, Spearman Correlation Coefficient | A2 |
| [48, 105] | Structure of Prediction Models | A2 |
| [37, 39] | Rank Correlation | A2, A4 |
| [26, 106] | Domain Experts Feedback | A3 |
| [67] | Error Probability | A3 |
| [55] | Global Confidence, Neighbors Density Confidence, Neighbors Similarity Confidence | A3 |
| [110] | LT15, LT30 | A3 |
| [61, 63] | Mean Rank Difference | A3 |
| [60] | Median Magnitude of the Relative Error (MdmRE) | A3 |
| [121] | Pareto Prediction Error | A3 |
| [10, 78, 96] | Rank Accuracy (RA) | A3 |
| [18, 67, 106, 114] | Tuning Time | A3 |
| [22, 60, 79] | Mean Square Error (MSE) | A3, A4 |
| [79] | p-value, R2, Residual Standard Error (RSE) | A4 |
| [14] | Reward | A4 |
| [73] | Statistical Significance | A4 |
| [73, 97] | Qualitative Analysis | A4, A5 |
| [49] | Ranking Constraints | A4, A5 |
| [77] | Delaney's Statistics | A5 |
| [77] | Distance Function, Hyper-volume (HV) | A5 |
| [24] | Equivalence among Combinatorial Models | A5 |
| [24] | Failure Index Delta (FID), Totally Repaired Models (TRM) | A5 |

Table 7. Validation procedure reported in the literature. A1: Pure Prediction; A2: Interpretability of Configurable Systems; A3: Optimization; A4: Dynamic Configuration; A5: Mining Constraints; A6: SPL Evolution.

source and target systems. Moreover, Kolesnikov et al. [48] and Valov et al. [105] use size metrics as insights of interpretability. Valov et al. [105] compare the structure (size) of prediction models built for the same configurable system and trained on different hardware platforms. They measured the size of a performance model by counting the features used in the nodes of a regression tree, while Kolesnikov et al. [48] define the model size metric as the number of configuration options in every term of the linear regression model. Temple et al. [100] reported that constraints extracted from decision trees were consistent with the domain knowledge of a video generator and could help developers preventing non-acceptable configurations. In their context, interpretability is important both for validating the insights of the learning and for encoding the knowledge into a variability model.

Sampling cost. Several works [6, 62, 63, 80, 109] use a sampling cost measurement to evaluate their prediction approach. In order to reduce measurement effort, these approaches aim at sampling configurations in a smart way by using as few configurations as possible. In [6, 62, 63, 109], sampling cost is considered as the total number of measurements required to train the model. These authors show the trade-off between the size of the training set and the accuracy of the model. Nair et al. [62, 63] compare different learning algorithms, along with different sampling techniques to determine exactly those configurations for measurement that reveal key performance characteristics and consequently reach the minimum possible sampling cost. Sarkar et al. [80] introduce a cost model, which they consider the measurement cost involved in building the training and testing sets, as well as the cost of building the prediction model and computing the prediction error. Sampling cost can be seen as an additional objective of the problem of learning configuration spaces, *i.e.* not only the usual learning accuracy should be considered.

The evaluation of the learning process requires to confront a trained model with new, un-measured configurations. Hold-out is the most considered strategy for splitting configuration data into training set and testing set. Depending on the targeted task (regression, ranking, classification, understanding), several metrics have been reused or defined. In addition to learning accuracy, sampling cost and interpretability of the learning model can be considered as part of the problem and thus assessed.

4.6 RQ6: What are the limitations faced by the current techniques and open challenges that need attention in the future?

In this section, we aim at identifying new areas of research that can lead to further enrichment of this field. Our recommendations are build upon previous insights and the analysis of papers' survey.

We start with the revisit of RQ1, RQ2, RQ3, RQ4, and RQ5.

Application objective. Despite a wide applicability being in terms of application domains, subject systems, or measurement properties, there is few concrete evidence that the proposed learning techniques are adopted in practice and daily used. To the best of our knowledge there is still a *gap*: we observe that end-users or developers of many configurable systems (*e.g.*, x264) are not yet using learning-based solutions. The reasons that prevent practical adoption are a first important issue to explore and understand. There are also research directions to further investigate to improve state-of-the-art approaches (see hereafter).

More applications. We report the use of learning techniques to target six main scenarios. However, there is still room to target other applications (*e.g.*, learning of multiple and composed configurable systems) or to combine scenarios (*e.g.*, mining constraint to support dynamic adaptation of multiple SPLs).

Integrated tooling support. We observe that several studies provide either open source or publicly available implementations. Tools could assist end-users during the configuration process and the choice of options. Variability management tools could be extended to encompass the learning stages reported in this SLR. It could also benefit to software developers in charge of maintaining configurable systems. However, none of the reported studies provide a tooling support fully integrated to existing SPL platforms or mainstream software tools. An exception is SPLConqueror [88] that supports an interactive configuration process.

Sampling. We reported 23 high-level sampling techniques used by state-of-the-art approaches. Most of the reported techniques only use as input a model of the configuration space (e.g., a feature model). The incorporation of knowledge and other types of input data in the sampling phase needs more investigation in future research (e.g., system documentation, implementation artifacts, code smells, and system evolution).

Constrained sampling. A difficult challenge when sampling is to generate configurations conform to constraints. It boils down to solve a satisfiability (SAT) problem. Recent results show that uniform, random sampling is still challenging for configurable systems [70]. Hence true random sampling is hard to instrument. In the context of testing SPL, there are numerous available heuristics, algorithms, and empirical studies [19, 50, 54, 56, 102]. In this case, sampling is used to efficiently test a subset of configurations. Though the sample has not the same purpose, we believe learning-based process could benefit from works done in the software testing domain. An interesting question is whether heuristics or sampling criterion devised for testing are as efficient for learning configuration spaces.

Sampling numeric options. The majority of sampling techniques operates over Boolean options. However, many configurable systems provide numerical options. In the literature, three numeric-sampling strategies have been reported: *random* [97, 98, 100], *Cartesian product of numerical sampling* [7], and *Plackett-Burman design* [84]. Numeric sampling have substantially different value ranges in comparison with binary sampling. The number of options' values to select can be huge while constraints should still be respected. Sampling with numeric options is still an open issue – not a pure SAT problem.

Measuring. There are a large number of NFPs reported in the literature (see Fig. 6 and Table 4), most of them being quantitative and performance properties. The measurement of performance is subject to intensive research: there can be bias, noise or uncertainty [5, 16, 103]. Quantifying the influence of such factors on the system behaviour is important to be able to avoid unexpected issues. Trubiani and Apel introduce the concept of uncertainty-influence models and present PLUS, a preliminary approach to learn and quantify the influence of uncertain parameters on system performance. Only a few works deal with uncertainty when learning configuration spaces (e.g., [41]). There are questions we could ask, such as how many times do we need to repeat measurements? Can we trust? Furthermore, there is a tension between the accuracy of measurements and the need to scale up the process to thousands of configurations. For example, the use of a cloud has the merit of offering numerous computing resources at the expense of networking issues and heterogeneity [10, 51]. Though some countermeasures can be considered, it remains an important practical problem. At the opposite end of the spectrum, simulation and static analysis are less costly but may approximate the real measures: few studies explore them which demand further investigation. Another issue is how performance measurements transfer to computational environment (e.g., hardware). The recent rise of transfer learning techniques is an interesting direction, but much more research is needed.

Learning. Numerous learning techniques have been proposed and can be applied: which one practitioners should use and under which conditions? It is still missing an actionable framework that would automatically choose or recommend a suited learning based on an engineering context (e.g., targeted NFPs, need of interpretability). Our SLR is a first step in this direction, since we

provide a repository of learning techniques together with their application objective. However more research is needed.

Parameter tuning and resampling. Most of the used learning techniques have parameters that guide their execution and have a strong influence on their prediction accuracy. However, most of the proposed approaches set these parameters manually based on the knowledge of a domain expert or even use default parameters. Also, the use of different resampling techniques is hardly explored by current works (except by [29, 38, 61]).

Unsupervised learning. Although all studies focused only on supervised learning, unsupervised learning is another potential candidate to support the exploration of large software configuration spaces. Still, analyzing the feasibility of unsupervised learning techniques remains unexplored.

Transfer learning. We notice that several recent works aim to transfer the learning-based knowledge of a configuration space to another setting (e.g., the use of another hardware). It is an important research direction, since many factors can influence software configuration spaces.

Validation. There is a wide range of validation metrics reported in the literature. We observe that some metrics can be used for the same task and there is not necessarily a consensus. On the other hand, most of the studies use an unique metric for validation, which may provide incomplete quantification of the accuracy.

Interpretability. There are few studies reported in the literature whose application objective is the comprehension of configurable systems. However, none of these studies use an empirical user studies to validate the proposed approach level of comprehension by domain experts. The existing approaches use only a model size metric to report the level of comprehension [48, 105].

Comprehensive evaluation. The reported studies analyse prediction error, model size, and computation time. As future work, we aim at analyzing to what extent the trade-offs among "prediction error, model size, and computation time" affect the application objective of prediction models: interactions, code smells, variability model characteristics (feature type, model height, etc), system evolution, and others. In addition, sampling cost could also be considered as part of the evaluation. Overall, we call for more comprehensive evaluation that aims to trade different aspects of concrete engineering context.

Unified Benchmark. Although we noticed the use of the same system across some studies, we also noticed the discrepancy of measurements among them. As future work, we aim at building a single trust benchmark repository with a portfolio of all data from the state-of-the-art systems in Table 4 and we advise researchers to use these benchmarks. Providing a repository of the measured subject systems from various domains will facilitate the comparison of different approaches under the same scenario.

Synergies with other community. As a final note, we notice that the problem of learning configuration spaces is at the intersection of:

- artificial intelligence (mainly statistical machine learning and constraint reasoning). Though many works in this field do not specifically address software systems or configuration problems [36, 74], there is an important potential to connect the dots and to reuse state-of-the-art methods for learning highly dimensional spaces like software configuration spaces;
- software engineering in general (including performance engineering and software testing). There are many challenges related to pure software: how to deploy, execute, and observe software configurations? How to synthesize knowledge actionable by software developers? Hence software frames the overall problem and calls to investigate specific issues (e.g., interpretability or transferability).

5 THREATS TO VALIDITY

This section discusses potential threats to validity that might have affected the results of the SLR. We faced similar threats to validity as any other SLR. The findings of this SLR may have been affected by bias in the selection of the primary studies, inaccuracy in the data extraction and in the classification of the primary studies, and incompleteness in defining the open challenges. Next, we summarize the main threats to the validity of our work and the strategies we have followed to minimize their impact. We discussed the SLR validity with respect to the two groups of common threats to validity: *internal* and *external* validity [112].

Internal validity. An internal validity threat concerns the reliability of the selection and data extraction process. To further increase the internal validity of the review results, the search for relevant studies was conducted in several relevant scientific databases, and it was focused not only on journals but also on conferences, symposiums, and workshops. Moreover, we conducted the inclusion and exclusion processes in parallel by involving three researchers and we cross-checked the outcome after each phase. In the case of disagreements, we discussed until a final decision was achieved. Furthermore, we documented potentially relevant studies that were excluded. Therefore, we believe that we have not omitted any relevant study. However, the quality of the search engines could have influenced the completeness of the identified primary studies (*i.e.*, our search may have missed those studies whose authors did not use the terms we used in our search string to specify keywords).

For the selected papers, a potential threat to validity is the reliability and accuracy of the data extraction process, since not all information was obvious to extract (*e.g.*, many papers lacked details about the measurement procedure and the validation design of the reported study). Consequently, some data had to be interpreted which involved subjective decisions by the researchers. Therefore, to ensure the validity, multiple sources of data were analyzed, *i.e.*, papers, websites, technical reports, manuals, and executable. Moreover, whenever there was a doubt about some extracted data in a particular paper, we discussed the reported data from different perspectives in order to resolve all discrepancies. However, we are aware that the data extraction process is a subjective activity and likely to yield different results when executed by different researchers.

External validity. A major external validity to this study was during the identified primary studies. Key terms are directly related to the *scope* of the paper and they can suffer a high variation. We limited the search for studies mainly targeting software systems (*e.g.*, software product lines) and thus mainly focus on software engineering conferences. This may affect the completeness of our search results since we are aware of some studies outside the software engineering community that also address the learning of software configurable systems. To minimize this limitation and avoid missing relevant papers, we also analyzed the references of the primary studies to identify other relevant studies. In addition, this SLR was based on a strict protocol described in Section 2 which was discussed before the start of the review to increase the reliability of the selection and data extraction processes of the primary studies and allow other researchers to replicate this review.

Another external validity concerns the description of open challenges. We are aware that the completeness of open challenges is another limitation that should be considered while interpreting the results of this review. It is also important to explore general contributions from other fields outside the software domain to fully gather the spread knowledge, which may extend the list of findings in this field. Therefore, in a future work, the list of findings highlighted in Section 4.6 may be extended by conducting an additional review, making use of other keywords to be able to find additional relevant studies outside the software community.

6 RELATED WORK

After the introduction of SLR in software engineering in 2004, the number of published reviews in this field has grown significantly [45]. A broad SLR has been conducted by Heradio et al. [34] to identify the most influential researched topics in SPL, and how the interest in those topics has evolved over the years. Although these reviews are not directly related to ours, the high level of detail of their research methodology supported to structure and define our own methodology.

Benavides et al. [12] presented the results of a literature review to identify a set of operations and techniques that provide support to the automatic analysis of variability models. In a similar scenario, Lisboa et al. [53] and Pereira et al. [68] conducted an SLR on variability management of SPLs. They reported several dozens of approaches and tools to support stakeholders in an interactive and automatic configuration process. Strict to the application engineering phase, Ochoa et al. [64] and Harman et al. [33] conducted a literature review on the use of search-based software engineering techniques for optimization of SPL configurations. In contrast to these previous reviews, our SLR provides further details on the use of automatic learning techniques to explore large configuration spaces. We contribute with a catalogue of sampling approaches, measurement procedures, learning techniques, and validation steps that serves as a summarization of the results in this specific field.

In [81], 14 software engineering experts and 229 Java software engineers were interviewed to identify major activities and challenges related to configuration engineering. In complement, a systematic literature review was conducted. In order to select papers, authors focused on those in which the title and abstract contain the keyword "config*", "misconfig*", or "mis-config*". On the one hand, the scope is much broader: It spans all engineering activities of configurable systems whereas our survey specifically targets learning-based approaches and applications. On the other hand, we learn that configuration spaces are studied in many different engineering contexts and we take care of considering a broader terminology (see Table 1, page 6). Despite different scope and reviewing methodology, the two surveys are complementary. A research direction is to further explore how learning-based approaches classified in this survey could support configuration engineering activities identified in [81].

Finally, complementary to our review, several works [19, 50, 54, 56, 102, 107] discussed product sampling techniques in the context of testing and verifying a software product line. They classify the proposed techniques into different categories and discuss the required input and criteria used to evaluate these techniques, such as the sample size, the rate of fault detection, and the tool support. Our SLR could benefit from their results through the use of sampling techniques still not explored by learning techniques (*e.g.* code and requirements coverage), as well as the SPL testing community could benefit from the sampling techniques reported in this paper still not used for testing. In [44], 111 real-world Web configurators are analyzed but do not consider learning mechanisms.

Overall, the main topics of previous SLRs include variability model analysis, variability management, configuration engineering, and SPL testing. None of the aforementioned surveys directly address the use of learning techniques to explore the behaviour of large configuration spaces.

7 CONCLUSION

We presented a systematic literature review related to the use of learning techniques to analyze large configuration software spaces. We analysed the literature in terms of a four-stage process: sampling, measuring, learning, and validation (see Section 3). Our contributions are fourfold. First, we identified the application of each approach which can guide researchers and industrial practitioners when searching for an appropriate technique that fits their current needs. Second, we classified the literature with respect to each learning stage. Mainly, we give an in-depth view of (i)

sampling techniques and employed design; (ii) measurement properties and effort for measurement; (iii) employed learning techniques; and (iv) how these techniques are empirically validated. Third, we provide a repository of all available subject systems used in the literature together with their application domains and qualitative or quantitative properties of interest. We welcome any contribution by the community: The list of selected studies and their classification can be found in our Web supplementary material [1]. Fourth, we identify the main shortcomings from existing approaches and non-addressed research areas to be explored by future work.

Our results reveal that the research in this field is application-dependent: Though the overall scheme remains the same ("sampling, measuring, learning"), the concrete choice of techniques should trade various criterion like safety, cost, accuracy, and interpretability. The proposed techniques have typically been validated with respect to different metrics depending on their tasks (e.g., performance prediction). Although the results are quite accurate, there is still need to decrease learning errors or to generalize predictions to multiple computing environments. Given the increasing interest and importance of this field, there are many exciting opportunities of research at the interplay of artificial intelligence and software engineering.

ACKNOWLEDGMENTS

This research was partially funded by the ANR-17-CE25-0010-01 VaryVary project. We would like to thank Paul Temple for his early comments on a draft of this article.

REFERENCES

- [1] 2019. Learning Software Configuration Spaces: A Systematic Literature Review. <https://github.com/VaryVary/ML-configurable-SLR> Accessed: 2019-06-04.
- [2] Mathieu Acher, Philippe Collet, Philippe Lahire, and Robert B France. 2013. Familiar: A domain-specific language for large scale management of feature models. *Science of Computer Programming (SCP)* 78, 6 (2013), 657–681.
- [3] 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.
- [4] Sheldon B. Akers. 1978. Binary decision diagrams. *IEEE Transactions on computers* 6 (1978), 509–516.
- [5] Aldeida Aleti, Catia Trubiani, André van Hoorn, and Pooyan Jamshidi. 2018. An efficient method for uncertainty propagation in robust software performance estimation. *Journal of Systems and Software* 138 (2018), 222–235. <https://doi.org/10.1016/j.jss.2018.01.010>
- [6] Omid Alipourfard, Hongqiang Harry Liu, Jianshu Chen, Shivaram Venkataraman, Minlan Yu, and Ming Zhang. 2017. Cherrypick: Adaptively unearthing the best cloud configurations for big data analytics. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*. 469–482.
- [7] 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.
- [8] Sven Apel, Don Batory, Christian Kästner, and Gunter Saake. 2013. *Feature-Oriented Software Product Lines: Concepts and Implementation*. Springer-Verlag.
- [9] Kacper Bąk, Zinovy Diskin, Michał Antkiewicz, Krzysztof Czarnecki, and Andrzej Wąsowski. 2016. Clafer: unifying class and feature modeling. *Software & Systems Modeling* 15, 3 (2016), 811–845.
- [10] Liang Bao, Xin Liu, Ziheng Xu, and Baoyin Fang. 2018. AutoConfig: automatic configuration tuning for distributed message systems. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. ACM, 29–40.
- [11] David Benavides, Pablo Trinidad Martín-Arroyo, and Antonio Ruiz Cortés. 2005. Automated reasoning on feature models.. In *International Conference on Advanced Information Systems Engineering (CAiSE)*, Vol. 5. Springer, 491–503.
- [12] David Benavides, Sergio Segura, and Antonio Ruiz-Cortés. 2010. Automated analysis of feature models 20 years later: a literature review. *Information Systems* 35, 6 (2010), 615–708.
- [13] Mikaela Cashman, Myra B. Cohen, Priya Ranjan, and Robert W. Cottingham. 2018. Navigating the maze: the impact of configurability in bioinformatics software. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 757–767. <https://doi.org/10.1145/3238147.3240466>
- [14] Haifeng Chen, Guofei Jiang, Hui Zhang, and Kenji Yoshihira. 2009. Boosting the performance of computing systems through adaptive configuration tuning. In *ACM Symposium on Applied Computing (SAC)*. ACM, 1045–1049.

- [15] Shipping 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.
- [16] Maxime Colmant, Romain Rouvoy, Mascha Kurpicz, Anita Sobe, Pascal Felber, and Lionel Seinturier. 2018. The next 700 CPU power models. *Journal of Systems and Software* 144 (2018), 382–396. <https://doi.org/10.1016/j.jss.2018.07.001>
- [17] Marco Couto, Paulo Borba, Jácome Cunha, João Paulo Fernandes, Rui Pereira, and João Saraiva. 2017. Products go green: Worst-case energy consumption in software product lines. In *Proceedings of the 21st International Systems and Software Product Line Conference-Volume A*. ACM, 84–93.
- [18] Yufei Ding, Jason Ansel, Kalyan Veeramachaneni, Xipeng Shen, Una-May O’Reilly, and Saman Amarasinghe. 2015. Autotuning algorithmic choice for input sensitivity. In *ACM SIGPLAN Notices*, Vol. 50. ACM, 379–390.
- [19] Ivan do Carmo Machado, John D Mcgregor, Yguaratã Cerqueira Cavalcanti, and Eduardo Santana De Almeida. 2014. On strategies for testing software product lines: A systematic literature review. *Information and Software Technology* 56, 10 (2014), 1183–1199.
- [20] Francisco Duarte, Richard Gil, Paolo Romano, Antónia Lopes, and Luís Rodrigues. 2018. Learning non-deterministic impact models for adaptation. In *Proceedings of the 13th International Conference on Software Engineering for Adaptive and Self-Managing Systems*. ACM, 196–205.
- [21] Holger Eichelberger, Cui Qin, Roman Sizonenko, and Klaus Schmid. 2016. Using IVML to model the topology of big data processing pipelines. In *International Systems and Software Product Line Conference (SPLC)*. ACM, 204–208.
- [22] Abdellatif El Afia and Malek Sarhani. 2017. Performance prediction using support vector machine for the configuration of optimization algorithms. In *2017 3rd International Conference of Cloud Computing Technologies and Applications (CloudTech)*. IEEE, 1–7.
- [23] Leire Etxeberria, Catia Trubiani, Vittorio Cortellessa, and Goiuria Sagardui. 2014. Performance-based selection of software and hardware features under parameter uncertainty. In *Proceedings of the 10th international ACM Sigsoft conference on Quality of software architectures*. ACM, 23–32.
- [24] Angelo Gargantini, Justyna Petke, and Marco Radavelli. 2017. Combinatorial interaction testing for automated constraint repair. In *2017 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 239–248.
- [25] Salah Ghamizi, Maxime Cordy, Mike Papadakis, and Yves Le Traon. 2019. Automated Search for Configurations of Deep Neural Network Architectures. *arXiv preprint arXiv:1904.04612* (2019).
- [26] 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.
- [27] Jianmei Guo, Krzysztof Czarnecki, Sven Apel, Norbert Siegmund, and Andrzej Wasowski. 2013. Variability-aware performance prediction: A statistical learning approach. In *ASE*.
- [28] Jianmei Guo, Jules White, Guangxin Wang, Jian Li, and Yinglin Wang. 2011. A genetic algorithm for optimized feature selection with resource constraints in software product lines. *Journal of Systems and Software* 84, 12 (2011), 2208–2221.
- [29] Jianmei Guo, Dingyu Yang, Norbert Siegmund, Sven Apel, Atrisha Sarkar, Pavel Valov, Krzysztof Czarnecki, Andrzej Wasowski, and Huiqun Yu. 2017. Data-efficient performance learning for configurable systems. *Empirical Software Engineering* (2017), 1–42.
- [30] 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* (July 2018). <https://doi.org/10.07980> Empirical Software Engineering journal.
- [31] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. 2009. The WEKA data mining software: an update. *ACM SIGKDD explorations newsletter* 11, 1 (2009), 10–18.
- [32] Svein O. Hallsteinsen, Mike Hinchey, Sooyong Park, and Klaus Schmid. 2008. Dynamic Software Product Lines. *IEEE Computer* 41, 4 (2008), 93–95. <https://doi.org/10.1109/MC.2008.123>
- [33] Mark Harman, Yue Jia, Jens Krinke, William B Langdon, Justyna Petke, and Yuanyuan Zhang. 2014. Search based software engineering for software product line engineering: a survey and directions for future work. In *Proceedings of the 18th International Software Product Line Conference-Volume 1*. ACM, 5–18.
- [34] Ruben Heradio, Hector Perez-Morago, David Fernandez-Amoros, Francisco Javier Cabrerizo, and Enrique Herrera-Viedma. 2016. A bibliometric analysis of 20 years of research on software product lines. *Information and Software Technology* 72 (2016), 1–15.
- [35] Rashina Hoda, Norsaremah Salleh, John Grundy, and Hui Mien Tee. 2017. Systematic literature reviews in agile software development: A tertiary study. *Information and software technology* 85 (2017), 60–70.
- [36] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2011. Sequential model-based optimization for general algorithm configuration. In *International Conference on Learning and Intelligent Optimization*. Springer, 507–523.

- [37] 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).
- [38] Pooyan Jamshidi and Giuliano Casale. 2016. An uncertainty-aware approach to optimal configuration of stream processing systems. In *2016 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*. IEEE, 39–48.
- [39] 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. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE Press, 497–508. <http://dl.acm.org/citation.cfm?id=3155625>
- [40] 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.
- [41] 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. <https://doi.org/10.1109/SEAMS.2017.11>
- [42] 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.
- [43] K. Kang, S. Cohen, J. Hess, W. Novak, and S. Peterson. 1990. *Feature-Oriented Domain Analysis (FODA)*. Technical Report CMU/SEI-90-TR-21. SEI.
- [44] Ebrahim Khalil Abbasi, Arnaud Hubaux, Mathieu Acher, Quentin Boucher, and Patrick Heymans. 2013. The Anatomy of a Sales Configurator: An Empirical Study of 111 Cases. In *CAiSE'13*.
- [45] Barbara Kitchenham, O Pearl Brereton, David Budgen, Mark Turner, John Bailey, and Stephen Linkman. 2009. Systematic literature reviews in software engineering—a systematic literature review. *Information and Software Technology (IST)* 51, 1 (2009), 7–15.
- [46] Barbara Kitchenham and Stuart Charters. 2007. Guidelines for performing Systematic Literature Reviews in Software Engineering. Citeseer.
- [47] 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).
- [48] Sergiy Kolesnikov, Norbert Siegmund, Christian Kästner, Alexander Grebhahn, and Sven Apel. 2018. Tradeoffs in modeling performance of highly configurable software systems. *Software & Systems Modeling* (2018), 1–19.
- [49] Thomas Krismayer, Rick Rabiser, and Paul Grünbacher. 2017. Mining constraints for event-based monitoring in systems of systems. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE Press, 826–831.
- [50] Jihyun Lee, Sungwon Kang, and Danhyung Lee. 2012. A survey on software product line testing. In *Proceedings of the 16th International Software Product Line Conference-Volume 1*. ACM, 31–40.
- [51] Philipp Leitner and Jürgen Cito. 2016. Patterns in the Chaos - A Study of Performance Variation and Predictability in Public IaaS Clouds. *ACM Trans. Internet Techn.* 16, 3 (2016), 15:1–15:23. <https://doi.org/10.1145/2885497>
- [52] 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).
- [53] Liana Barachisio Lisboa, Vinicius Cardoso Garcia, Daniel Lucrédio, Eduardo Santana de Almeida, Silvio Romero de Lemos Meira, and Renata Pontin de Mattos Fortes. 2010. A systematic review of domain analysis tools. *Information and Software Technology* 52, 1 (2010), 1–13.
- [54] Roberto E Lopez-Herrejon, Stefan Fischer, Rudolf Ramler, and Alexander Egyed. 2015. A first systematic mapping study on combinatorial interaction testing for software product lines. In *2015 IEEE Eighth International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 1–10.
- [55] Jabier Martinez, Jean-Sébastien Sottet, Alfonso García Frey, Tegawendé F Bissyandé, Tewfik Ziadi, Jacques Klein, Paul Temple, Mathieu Acher, and Yves Le Traon. 2018. Towards Estimating and Predicting User Perception on Software Product Variants. In *International Conference on Software Reuse*. Springer, 23–40.
- [56] 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*. ACM, 643–654.
- [57] 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.

- [58] Brice Morin, Olivier Barais, Jean-Marc Jézéquel, Franck Fleurey, and Arnor Solberg. 2009. Models@ Run.time to Support Dynamic Adaptation. *IEEE Computer* 42, 10 (2009), 44–51. <https://doi.org/10.1109/MC.2009.327>
- [59] Alexandr Murashkin, Michał Antkiewicz, Derek Rayside, and Krzysztof Czarnecki. 2013. Visualization and exploration of optimal variants in product line engineering. In *Proceedings of the 17th International Software Product Line Conference*. ACM, 111–115.
- [60] 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>
- [61] 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>
- [62] 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.
- [63] Vivek Nair, Zhe Yu, Tim Menzies, Norbert Siegmund, and Sven Apel. 2018. Finding faster configurations using flash. *IEEE Transactions on Software Engineering* (2018).
- [64] Lina Ochoa, Oscar Gonzalez-Rojas, Alves Pereira Juliana, Harold Castro, and Gunter Saake. 2018. A systematic literature review on the semi-automatic configuration of extended product lines. *Journal of Systems and Software* 144 (2018), 511–532.
- [65] Lina Ochoa, Oscar González-Rojas, and Thomas Thüm. 2015. Using decision rules for solving conflicts in extended feature models. In *International Conference on Software Language Engineering (SLE)*. ACM, 149–160.
- [66] 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. <https://doi.org/10.1145/3106237.3106273>
- [67] Takayuki Osogami and Sei Kato. 2007. Optimizing system configurations quickly by guessing at the performance. In *ACM SIGMETRICS Performance Evaluation Review*, Vol. 35. ACM, 145–156.
- [68] Juliana Alves Pereira, Kattiana Constantino, and Eduardo Figueiredo. 2015. A systematic literature review of software product line management tools. In *International Conference on Software Reuse (ICSR)*. Springer, 73–89.
- [69] Juliana Alves Pereira, Pawel Matuszyk, Sebastian Krieter, Myra Spiliopoulou, and Gunter Saake. 2018. Personalized recommender systems for product-line configuration processes. *Computer Languages, Systems & Structures* (2018).
- [70] 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>
- [71] Klaus Pohl, Günter Böckle, and Frank J. van der Linden. 2005. *Software product line engineering: foundations, principles and techniques*. Springer, Berlin Heidelberg.
- [72] Richard Pohl, Kim Lauenroth, and Klaus Pohl. 2011. A performance comparison of contemporary algorithmic approaches for automated analysis operations on feature models. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE Computer Society, 313–322.
- [73] Adam Porter, Cemal Yilmaz, Atif M Memon, Douglas C Schmidt, and Bala Natarajan. 2007. Skoll: A process and infrastructure for distributed continuous quality assurance. *IEEE Transactions on Software Engineering* 33, 8 (2007), 510–525.
- [74] Suknawati Anggraeni Putri et al. 2017. Combining integrated sampling technique with feature selection for software defect prediction. In *2017 5th International Conference on Cyber and IT Service Management (CITSM)*. IEEE, 1–6.
- [75] 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.
- [76] Fabricia Roos-Frantz, David Benavides, Antonio Ruiz-Cortés, André Heuer, and Kim Lauenroth. 2012. Quality-aware analysis in product line engineering with the orthogonal variability model. *Software Quality Journal* 20, 3-4 (2012), 519–565.
- [77] Safdar Aqeel Safdar, Hong Lu, Tao Yue, and Shaukat Ali. 2017. Mining cross product line rules with multi-objective search and machine learning. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 1319–1326.
- [78] Muhammad Suleman Saleem, Chen Ding, Xumin Liu, and Chi-Hung Chi. 2015. Personalized decision-strategy based web service selection using a learning-to-rank algorithm. *IEEE Transactions on Services Computing* 8, 5 (2015), 727–739.
- [79] 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.
- [80] Atri Sarkar, Jianmei Guo, Norbert Siegmund, Sven Apel, and Krzysztof Czarnecki. 2015. Cost-efficient sampling for performance prediction of configurable systems (t). In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 342–352.

- [81] Mohammed Sayagh, Nouredine Kerzazi, Bram Adams, and Fabio Petrillo. 2018. Software Configuration Engineering in Practice: Interviews, Survey, and Systematic Literature Review. *IEEE Transactions on Software Engineering* (2018).
- [82] Abdel Salam Sayyad, Tim Menzies, and Hany Ammar. 2013. On the value of user preferences in search-based software engineering: a case study in software product lines. In *Proceedings of the 2013 International Conference on Software Engineering*. IEEE Press, 492–501.
- [83] Amir Molzam Sharifloo, Andreas Metzger, Clément Quinton, Luciano Baresi, and Klaus Pohl. 2016. Learning and evolution in dynamic software product lines. In *2016 IEEE/ACM 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*. IEEE, 158–164.
- [84] Norbert Siegmund, Alexander Grebhahn, Sven Apel, and Christian Kästner. 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.
- [85] 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.
- [86] Norbert Siegmund, Marko Rosenmüller, Christian Kästner, Paolo G. Giarrusso, Sven Apel, and Sergiy S. Kolesnikov. 2011. Scalable prediction of non-functional properties in software product lines. In *Software Product Line Conference (SPLC), 2011 15th International*. 160–169.
- [87] Norbert Siegmund, Marko Rosenmüller, Christian Kästner, Paolo G. Giarrusso, Sven Apel, and Sergiy S. Kolesnikov. 2013. Scalable prediction of non-functional properties in software product lines: Footprint and memory consumption. *Information and Software Technology* 55, 3 (2013), 491–507.
- [88] Norbert Siegmund, Marko Rosenmüller, Martin Kuhlemann, Christian Kästner, Sven Apel, and Gunter Saake. 2012. SPL Conqueror: Toward optimization of non-functional properties in software product lines. *Software Quality Journal* 20, 3-4 (2012), 487–517.
- [89] 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.
- [90] Norbert Siegmund, Stefan Sobernig, and Sven Apel. 2017. Attributed variability models: outside the comfort zone. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. ACM, 268–278.
- [91] Norbert Siegmund, Alexander von Rhein, and Sven Apel. 2013. Family-based performance measurement. In *ACM SIGPLAN Notices*, Vol. 49. ACM, 95–104.
- [92] Julio Sincero, Wolfgang Schroder-Preikschat, and Olaf Spinczyk. 2010. Approaching non-functional properties of software product lines: Learning from products. In *Software Engineering Conference (APSEC), 2010 17th Asia Pacific*. 147–155.
- [93] 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.
- [94] Jeffrey Stuckman, James Walden, and Riccardo Scandariato. 2017. The effect of dimensionality reduction on software vulnerability prediction models. *IEEE Transactions on Reliability* 66, 1 (2017), 17–37.
- [95] Mikael Svahnberg, Jilles van Gorp, and Jan Bosch. 2005. A taxonomy of variability realization techniques: Research Articles. *Softw. Pract. Exper.* 35, 8 (2005), 705–754. <https://doi.org/10.1002/spe.v35:8>
- [96] Ivan Švogor, Ivica Crnković, and Neven Vrček. 2019. An extensible framework for software configuration optimization on heterogeneous computing systems: Time and energy case study. *Information and software technology* 105 (2019), 30–42.
- [97] 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).
- [98] 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>
- [99] Paul Temple, Mathieu Acher, Jean-Marc A. Jézéquel, Léo A. Noel-Baron, and José A. Galindo. 2017. *Learning-Based Performance Specialization of Configurable Systems*. Research Report. IRISA, Inria Rennes ; University of Rennes 1. <https://hal.archives-ouvertes.fr/hal-01467299>
- [100] 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://doi.org/10.1145/2934466.2934472>
- [101] 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.
- [102] Thomas Thüm, Sven Apel, Christian Kästner, Ina Schaefer, and Gunter Saake. 2014. A classification and survey of analysis strategies for software product lines. *ACM Computing Surveys (CSUR)* 47, 1 (2014), 6.

- [103] Catia Trubiani and Sven Apel. 2019. PLUS: Performance Learning for Uncertainty of Software. In *International Conference on Software Engineering NIER*. ACM.
- [104] Pavel Valov, Jianmei Guo, and Krzysztof Czarnecki. 2015. Empirical comparison of regression methods for variability-aware performance prediction. In *SPLC'15*.
- [105] 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.
- [106] Dana Van Aken, Andrew Pavlo, Geoffrey J Gordon, and Bohan Zhang. 2017. Automatic database management system tuning through large-scale machine learning. In *Proceedings of the 2017 ACM International Conference on Management of Data*. ACM, 1009–1024.
- [107] 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*. ACM, 1–13.
- [108] Sravanthi Kota Venkata, Ikkjin Ahn, Donghwan Jeon, Anshuman Gupta, Christopher Louie, Saturnino Garcia, Serge Belongie, and Michael Bedford Taylor. 2009. SD-VBS: The San Diego vision benchmark suite. In *2009 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 55–64.
- [109] Markus Weckesser, Roland Kluge, Martin Pfannemüller, 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.
- [110] Dennis Westermann, Jens Happe, Rouven Krebs, and Roozbeh Farahbod. 2012. Automated inference of goal-oriented performance prediction functions. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. ACM, 190–199.
- [111] Claes Wohlin. 2014. Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the 18th international conference on evaluation and assessment in software engineering*. ACM, 38.
- [112] C Wohlin, P Runeson, M Host, MC Ohlsson, B Regnell, and A Wesslen. 2000. Experimentation in software engineering: an introduction.
- [113] Bawei Xi, Zhen Liu, Mukund Raghavachari, Cathy H Xia, and Li Zhang. 2004. A smart hill-climbing algorithm for application server configuration. In *Proceedings of the 13th international conference on World Wide Web*. ACM, 287–296.
- [114] Lin Xu, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2008. SATzilla: portfolio-based algorithm selection for SAT. *Journal of artificial intelligence research* 32 (2008), 565–606.
- [115] Tianyin Xu, Long Jin, Xuepeng Fan, Yuanyuan Zhou, Shankar Pasupathy, and Rukma Talwadker. 2015. Hey, you have given me too many knobs!: understanding and dealing with over-designed configuration in system software. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015, Bergamo, Italy, August 30 - September 4, 2015*. 307–319. <https://doi.org/10.1145/2786805.2786852>
- [116] 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.
- [117] Cemal Yilmaz, Emine Dumlu, Myra B Cohen, and Adam Porter. 2014. Reducing masking effects in combinatorial-interaction testing: A feedback driven adaptive approach. *IEEE Transactions on Software Engineering* 40, 1 (2014), 43–66.
- [118] Yi Zhang, Jianmei Guo, Eric Blais, and Krzysztof Czarnecki. 2015. Performance prediction of configurable software systems by fourier learning (t). In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 365–373.
- [119] 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.
- [120] Wei Zheng, Ricardo Bianchini, and Thu D Nguyen. 2007. Automatic configuration of internet services. *ACM SIGOPS Operating Systems Review* 41, 3 (2007), 219–229.
- [121] Marcela Zuluaga, Andreas Krause, and Markus Püschel. 2016. ϵ -pal: an active learning approach to the multi-objective optimization problem. *The Journal of Machine Learning Research* 17, 1 (2016), 3619–3650.