

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

1-2022

Predictive models in software engineering: Challenges and opportunities

Yanming YANG

Xin XIA

David LO

Singapore Management University, davidlo@smu.edu.sg

Tingting BI

John C. GRUNDY

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#), and the [Software Engineering Commons](#)

Citation

YANG, Yanming; XIA, Xin; LO, David; BI, Tingting; GRUNDY, John C.; and YANG, Xiaohu. Predictive models in software engineering: Challenges and opportunities. (2022). *ACM Transactions on Software Engineering and Methodology*. 31, (3), 1-72.

Available at: https://ink.library.smu.edu.sg/sis_research/7630

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Yanming YANG, Xin XIA, David LO, Tingting BI, John C. GRUNDY, and Xiaohu YANG

Predictive Models in Software Engineering: Challenges and Opportunities

YANMING YANG, Zhejiang University, China, and Monash University, Australia

XIN XIA, Monash University Australia

DAVID LO, Singapore Management University, Singapore

TINGTING BI, Monash University, Australia

JOHN GRUNDY, Monash University, Australia

XIAOHU YANG, Zhejiang University, China

Predictive models are one of the most important techniques that are widely applied in many areas of software engineering. There have been a large number of primary studies that apply predictive models and that present well-preformed studies and well-designworks in various research domains, including software requirements, software design and development, testing and debugging and software maintenance. This paper is a first attempt to systematically organize knowledge in this area by surveying a body of 139 papers on predictive models. We describe the key models and approaches used, classify the different models, summarize the range of key application areas, and analyze research results. Based on our findings, we also propose a set of current challenges that still need to be addressed in future work and provide a proposed research road map for these opportunities.

Additional Key Words and Phrases: Predictive models, machine learning, software engineering, survey

ACM Reference format:

Yanming Yang, Xin Xia, David Lo, Tingting Bi, John Grundy, and Xiaohu Yang . 2016. Predictive Models in Software Engineering: Challenges and Opportunities. 1, 1, Article 1 (January 2016), 35 pages.

DOI: 10.1145/nnnnnnnn.nnnnnnn

1 INTRODUCTION

Researchers have developed automated methodologies to improve software engineering tasks. Key reasons are usually to save developer time and effort and to improve the software quality in terms of stability, reliability, and security. Many of such studies have resulted in great improvements in various tasks [9, 62, 116, 136, 158].

A key technology, the *predictive model*, has been developed to solve a range of software engineering problems over several decades. The use of predictive models is in fact becoming increasingly popular in a wide range of software engineering research areas. Predictive models are built based on different types of datasets – such as software requirements, APIs, bug reports, source code and run-time data – and provide a final output according to distinct features found in the data. There are various predictive models commonly used in software engineering tasks that contribute to improving the efficiency of development processes and software quality. Common ones include defect prediction [136], API issue classification [60], and code smell detection [102].

Despite numerous studies on predictive models in software engineering, to the best of our knowledge, there has been no systematic study to analyze the use of and demonstrated the potential value of and current challenges of using predictive models in software engineering. There is no clear answer as to which software engineering tasks predictive models can be best applied and how to best go about leveraging the right predictive models for

2016. XXXX-XXXX/2016/1-ART1 \$15.00

DOI: 10.1145/nnnnnnnn.nnnnnnn

, Vol. 1, No. 1, Article 1. Publication date: January 2016.

these tasks. Answering these questions would be beneficial for both practitioners and researchers, in order to make informed decisions to solve a problem or conduct research using predictive models.

This paper contributes to the research on predictive models by performing a comprehensive systematic survey of the domain. At first we recorded all titles of papers published in top software engineering venues, such as ICSE, ASE, FSE, TSE, TOSEM, and EMSE from 2009 – 2019. We performed searches on the title and abstract of each paper and collected over 2,000 papers containing our target search terms. We then carefully read the abstract and introduction of each paper to validate its relevance to this study. Finally, we identified 145 high-quality primary study papers. We found that predictive models have been widely applied to a variety of software engineering tasks. Common ones include requirements classification, code change detection and malware detection. We grouped predictive model usage in software engineering into six research domains according to the software development lifecycle (SDLC) – software requirements, software design and development, testing and debugging, software maintenance, and software professional practice knowledge area. We further summarized and analyzed these studies in each domain. After that, we identified a set of remaining research and practice challenges and new research directions that we think should be investigated in the future. To the best of our knowledge, we are the first to perform such a systematic review on the use of predictive models in software engineering domain.

This paper makes the following key contributions:

- We present a comprehensive survey on predictive models covering 145 primary study papers from the last decade;
- We analyze these 145 primary studies and characterize them, and identify 10 commonly used measures for model evaluation;
- We summarize the research domains in which the predictive model has been applied, and discuss distinct technical challenges of using predictive models in software engineering; and
- We outline key future avenues for research on predictive models in software engineering.

The remainder of this paper is organized as follows. Section 2 briefly introduces the workflow of predictive models. Section 3 presents our study methodology. Section 4 investigates the evolution and distribution of the selected primary studies using predictive models for software engineering tasks, and Section 5 gives a classification and research distribution of these predictive models. Section 6 summarizes and analyzes the application of predictive models in software engineering, and a discussion of the threats that could affect the validity of our findings is presented in Section 7. Section 8 discusses the challenges that still need to be addressed in future work and give outlines a research road map of potential research opportunities in this domain. Section 9 provides a summary of the key conclusions of this work.

2 PREDICTIVE MODELS

2.1 Workflow of predictive models

A predictive model is usually treated as a black box that automatically assigns a class label when presented with the feature set of an unknown record, illustrated in Fig. 1. Instances in the input dataset can be assigned to one of several predefined categories by building such a predictive model. The predictive model is usually trained with a representative input dataset, then applied to target input datasets. The workflow of using a predictive model can be described as a mathematical problem of learning a target function (f) that maps each feature set (x) in a dataset (X) to one class label (y). The job of predictive models is to find the best target function (f_m).

There are four critical components when building such a predictive model:

Datasets: As the most basic component of predictive models, the dataset has a large impact on a model's performance. Low-quality datasets with noise and mislabeling may lead predictive models to provide (very) wrong experimental results, even if the process of model selection and training is effective.

Different types of datasets are used when performing different software engineering tasks with predictive models. For instance, studies may use source code, bug reports, or requirement documents as datasets of predictive models for defect prediction, bug report classification, and requirements-related knowledge classification respectively. In general, different datasets have different properties. These include scale, distribution, bias, quality, representativeness, sparseness and so on.

Features: In building predictive models, features (or attributes) in datasets are essential in the model training phase. The goal of a predictive model is essentially to learn a target function by analyzing different feature sets in given input datasets. Thus a good feature set can allow predictive models to learn the potential patterns in datasets and thus to output correct labels effectively. In order to construct high-quality feature sets, feature selection is an important optimization technique in building predictive models. Better feature selection strategies can improve the accuracy of results, reduce overfitting, as well as reduce training time.




Fig. 1. Predictive model as the task of mapping an input feature set x into its class label y .

Model Building Algorithms: A predictive model can be implemented by using different algorithms. For instance, J48, C4.5 and CART are commonly used algorithms for building Decision Tree-based predictive models. A number of relatively new deep learning-based algorithms and architectures, such as RNN and CNN-based networks, have been applied as predictive models to many software engineering tasks. Some predictive model algorithms have been introduced for specific software engineering problems, by enhancing existing algorithms, or providing new model capabilities. For example, many large and complex software-intensive systems widely used logs for troubleshooting, and thus Zhang et al. [160] present a tool, called LogRobust, to detect system anomalies by analyzing log messages. They adopted Bi-LSTM, a variant of LSTM model, to capture the contextual information in log sequences, and combined attention mechanism to strengthen the ability of automatically learning the importance of different log events. Liu et al. [63] concentrated on detecting feature envy, one of the most common code smells, and tailored a deep learning model to consider both textual input (method name or class name) and numerical input (the distance between a method and a class) by merging two CNN models.

Model Performance Measures: Many widely-used evaluation metrics, such as recall, precision, F-measure, accuracy, and AUC, have been introduced and applied in software engineering studies. This is because different tasks may require different metrics to evaluate the effectiveness of their proposed predictive models. On the other hand, some of the less common evaluation metrics, such as Balance, G-measure, and specificity, also appear in a few studies. We will present an analysis of various evaluation metrics for predictive models in Section 5.

3 METHODOLOGY

To perform a comprehensive systematic review of the use of predictive models in software engineering, we followed the systematic review guidelines provided by Kitchenham and Charters [51] and Petersen et al. [109].

3.1 Research Questions

The aim of this paper is to summarize, classify, analyze and propose research directions based on empirical evidence concerning the different modeling techniques and task domains involved. To achieve this, we define three research questions as shown in Table 1 and give the motivation behind each question.

Table 1. Research Questions and Motivations

RQs	Research Question	Motivation
RQ1	What are the trends in the primary studies on use of predictive models in SE?	The basic information (e.g., authors, publication year, affiliations) of the primary study papers can be informative. We wanted to understand where primary study papers on SE uses of predictive models are being published and any trends that can be observed. The goal of this RQ is to investigate these publication trends and distribution of the primary studies.
RQ2	Which predictive models are applied to what software engineering tasks?	Various predictive models are used to in the context of different software engineering tasks. It is necessary to generalize about predictive models so that researchers can select the appropriate model when researching in certain domains. The goal of this question is to determine which predictive models are frequently applied in software engineering, and to develop a classification for them.
RQ3	In what software engineering domains and applications have predictive models been applied?	Although predictive models have been applied to broad application scenarios in software engineering, there does not exist any comprehensive study that summarizes these research domains and applications using different models. This RQ aims to find and report on such a domain analysis.

RQ1 will analyze the distribution of publications on predictive models over the last decade to give an overview of the trend in software engineering research. RQ2 will provide a classification and distribution of predictive models used in software engineering and identify what are the commonly used evaluation metrics. RQ3 will explore where and how predictive models have been applied for specific software engineering tasks.

3.2 Literature Search and Selection

We identified a set of search terms based on software engineering tasks involved in predictive models that were already known to us. We then refined these by checking the title and abstract of the relevant papers. After considering the alternative spellings and synonyms for these search terms [38], they were combined with logical ORs, forming the complete search string. The search terms are listed as follows:

("predict" OR "identif*" OR "detect*" OR "classify*" OR "infer*" OR "recommend*" OR "indicat*" OR "estimat*") Since the Digital Library can support search terms that are not case sensitive and not "whole words only", we use the stemmed term to represent all forms related to them (e.g., classify* can search for classify, classifying, classified as well as classification).*

Following previous survey study approaches [38, 41], we first collected the titles of all papers published at ICSE, ASE, FSE, TSE, TOSEM, and EMSE between 2009 and 2019 from DBLP computer science bibliography. We then used our search string to perform search for papers related to predictive models, downloaded relevant papers from IEEEExplore, ACM Digital Library, ISI Web of Science, and Google Scholar. We obtained 2,128 candidate primary studies; the initial number of studies for each search engine is shown in Table 2. After discarding duplicate papers, we applied inclusion/exclusion criteria (see Section 2.3) to the title, abstract and keywords of each paper for narrowing the candidate set size to 521 studies. We then validated their relevance by reading these studies in full and found that a final 145 fulfilled all criteria.

Table 2. The initial number of studies for different search engines

Search Engine	#Studies
ACM Digital Library	402
IEEE Xplore Digital Library	273
ISI Web of Science	628
Google Scholar	636
DBLP	189
Total	2128

3.3 Inclusion and Exclusion Criteria

A number of low quality studies (e.g., non-published manuscripts and grey literature) are found when searching for the relevant papers with search engines. We also wanted only the most comprehensive or latest version of repeated

studies in the candidate set. For example, we removed a conference paper if it was extended in a follow-on journal paper. Thus we applied a set of inclusion and exclusion criteria to select the studies with strong relevance.

The following inclusion and exclusion criteria were used:

- ✓ The paper must be written in English.
- ✓ The paper must apply one or more predictive models to address software engineering tasks.
- ✓ The paper must be a peer reviewed full research paper published in a conference proceedings or a journal.
- ✗ Papers less than six pages are not considered.
- ✗ The predictive models used in studies cannot be baseline approaches.
- ✗ Conference version of a study that has an extended journal version is not considered.
- ✗ Keynote records and publications in grey literature are dropped.

The inclusion and exclusion criteria were piloted and applied by the first and fourth authors by assessing 45 randomly selected papers from the initial set. We used Cohen's kappa coefficient to measure the reliability of the inclusion and exclusion decisions. The agreement rate in the pilot study was "substantial" (0.72). We also performed the assessment in the full list of identified papers, and Cohen's kappa score was "substantial" (0.68). The first and fourth authors reached a consensus through discussions when encountering disagreements. In case disagreements were not resolved, a third researcher who is not the author of the paper was invited to make the final decision.

3.4 Data Extraction and Collection

After carefully reading all 145 papers in full, we extracted the required data and conducted detailed analysis to answer our three research questions. The detailed information is summarized in Table 3. Data collection mainly concentrated on three aspects: the fundamental information of each paper, some contents about predictive models and the research domain to which each study belongs. In order to prevent data loss and avoid mistakes as much as possible, data collection was performed by the first and fourth authors and then the results were verified by other researchers who are not co-authors of this paper.

Table 3. Data Collection for Research Questions

RQs	Types of Data to be Extracted
RQ1	Fundamental information for each paper (publication year, author name, type of paper and affiliation).
RQ2	Description, analysis and performance of different predictive models.
RQ3	Background, motivation, application scenario, and research topic of each study.

4 RQ1: WHAT ARE THE TRENDS IN THE PRIMARY STUDIES ON USE OF PREDICTIVE MODELS IN SE?

To discuss the emerging trends in predictive model use in SE, in this section we present an analysis of primary studies from three perspectives according to the fundamental information presented in these studies.

4.1 Publication trends in research on predictive models in Software Engineering

Fig. 2(a) shows the number of papers published between January 1st, 2009 and December 31st, 2019. It is clear that the number of publications from 2009 to 2013 shows a marked increase, with the number reaching 16 publications in 2013. Furthermore, except for 2014 and 2017 with a relatively small number of publications, the number of publications has stayed at a high level and reaches 26 in 2019.

An analysis of the cumulative publications is shown in Fig. 2(b). We used a polynomial function to fit the cumulative number of publications, revealing the publication trend in the last decade. It can be observed that The




Fig. 2. Publication trends of predictive models in the last decade.

(R^2) amounts to 0.99803 and the slope of the curve increases markedly between 2009 and 2019. This indicates that research studies using predictive models are likely to continue to experience a strong growth in the future. According to the trend of this curve, it can be foreseen that the cumulative number of publications may be over 170 by the end of 2020. We can see from Fig. 2 that the use of predictive models is becoming more popular in software engineering and there will be more studies that use these models to tackle practical SE problems.

4.2 Distribution of Publication Venues and Types of Contributions

The 145 reviewed studies were published in various publication venues, including in 3 top journals and 3 top conferences, well-known and highly regarded in the field of software engineering. The prevalence of papers on predictive models in these conferences and journals indicates construction of predictive models for software engineering purposes are considered of importance. Table 4 lists the number of papers published in each publication venue. ICSE (47) includes the highest number of primary study papers compared with other top conferences and journals; almost two times higher than ASE (22). In Table 4, TSE is the most popular journal containing almost 19% of the primary studies, followed by EMSE and FSE (21). There are 7 relevant studies published in TOSEM. We can observe that the number of conference papers employing predictive models is higher than that of journal papers.

Table 4. Top Venues with a Minimum of Two Predictive Model Based Papers

Rank	Acronym	Full name	#Studies
1.	ICSE	ACM/IEEE International Conference on Software Engineering	47
2.	TSE	IEEE Transactions on Software Engineering	27
3.	ASE	IEEE/ACM International Conference Automated Software Engineering	22
4.	EMSE	Empirical Software Engineering (published by Springer)	21
5.	FSE	ACM SIGSOFT Symposium on the Foundation of Software Engineering/ European Software Engineering Conference	21
6.	TOSEM	ACM Transactions on Software Engineering and Methodology	7

We checked the distribution between different publication venues. We find that 62% of publications appeared in conferences while 38% were published in journals. Fig. 3 gives the venue distribution per year. As From 2009 to 2016, most of the primary studies were published in conference proceedings. This trend has however changed, with the number of journal papers now increasing between 2014 and 2018, often outnumbering conference papers.




Fig. 3. Venue distribution per year.




Fig. 4. Type of main contribution.

Fig. 4 categorizes the primary publications according to their main contribution. The main contribution of 41% publications was to propose a novel predictive model technique / methodology. 25% focused on assessment and empirical studies, and 24% were case studies. The main contribution of 9 studies was to propose a prototype tool (6%), and 4% of primary studies were survey or literature reviews on predictive models.

Summary of answers to RQ1:

- (1) 145 relevant papers were identified from six top publication venues. The predictive model has attracted sustained interest, with the topic showing steady increase in primary studies in 2009-2019.
- (2) Most studies were published at conferences, compared with journals, though the trend changes in recent years. And the main contribution of most studies was to present a new technique or methodology.

5 RQ2: WHICH PREDICTIVE MODELS ARE APPLIED TO SOFTWARE ENGINEERING TASKS?

We present an overview of predictive models, including a classification and research distribution of these models from the primary studies. We also identify several key evaluation metrics that are used in these primary studies for evaluating the performance of their predictive models.

5.1 Predictive Model Classification

The predictive model, as a widely-used technology, has many different implementations. Those modelling techniques can be classified into four categories based on [38]: (1) Base learners, e.g., naive Bayes, Logistic Regression, (2) Ensemble learning techniques, e.g., Boosting, Bagging, and Random Forest, (3) Neural network, e.g., MLP, Deep Belief Network, CNN, RNN, and LSTM.

Base learners: Base learners are usually generated from training data by a base learning algorithm [166]. These are commonly a decision tree, neural network (e.g., MLP, etc.) or other kind of machine learning algorithms. 84 out of 150 primary studies only used base learners to build predictive models. Most of them are classic algorithms which work well for many tasks.

Ensemble learning techniques: Most ensemble methods use a single base learning algorithm to produce homogeneous base learners. However there are also some methods which use multiple learning algorithms to produce heterogeneous learners for reducing bias (Boosting), variance (Bagging), or improving predictions. The well-known Random Forest algorithm, as a parallel ensemble algorithm, has become one of the most common used ensemble algorithms. A large number of studies (22) employed Random Forest to address different software engineering research tasks. These include defect prediction, vulnerability prediction, code quality prediction, software license exception detection, and fault localization [87, 105, 137, 142, 163].

There are 9 studies using a Boosting approach in their reported models ([22, 68, 69, 71, 99, 104, 150]). Most of these studies ([22, 71, 150]) utilized LogitBoost as a part of their approach. Yang et al. [150] implemented a prototype tool with LogitBoost to detecting nocuous coordination ambiguities in requirement documents, which pose a high risk of misunderstanding among different developers. Falessiet al. [22] introduced a novel approach to estimate the number of remaining positive links in traceability recovery with LogitBoost. There are some studies that use other Boosting techniques, such as RankBoost and Gradient Boost. Perini et al. [104] employed RankBoost in order to give priority to requirements that were to be considered first. Machalica et al. [68] present a data-driven test selection strategy by training a gradient boosted decision tree classifier.

Bagging techniques were also applied by 7 studies to make predictions [54, 164]. Zhou et al. [164] leveraged six representative machine learning methods including bagging, AdaBoost and Random Forest to build fault prediction models to study the potentially confounding effect of class size. Kim et al. [54] evaluated the impact of data noise on defect prediction with adoption of basic learners and bagging learners. Malhotra and Khanna [71] evaluated the performance of LogitBoost, AdaBoost, Bagging, and other ML techniques for handling imbalanced datasets in software change prediction task.

Neural Network: Neural networks (NN) are a set of algorithms (e.g., perceptions and deep learning techniques) designed to identify patterns. They interpret sensory data through a machine perception, labeling, or classifying raw data input. Therefore, NN can be considered as components of larger machine-learning applications involving algorithms for clustering, classification, or regression. Neural Network models can be classified into three categories, i.e., perceptron models, MP Neural models, and deep learning models, through the analysis of selected studies.

A perceptron is the simplified form of a neural network applied for supervised learning of binary classifiers, which consists of four main components including input values, net sum, weights and bias, and an activation function.

The MP neuron model is actually a simplified model constructed based on the structure and working principle of biological neurons, which consists of a function with a single parameter, taking binary input, and giving binary output according to a determined threshold value.

Deep learning is part of a broader family of machine learning methods based on artificial neural networks [121]. Deep learning architectures, such as deep neural networks, deep belief networks, recurrent neural networks (RNN), and convolutional neural networks (CNN), have been applied to help various tasks across the software development life cycle. With the advent and development of deep learning techniques, 15 primary studies utilized deep learning techniques in their reported models [10, 61, 63, 149]. In this section, we briefly introduce several deep neural networks to give an overview of the fundamental principles. We also present how they have been applied to solve software engineering problems.

CNN has been used for SE tasks, e.g., [63]. The CNN model consists of three layers, including an input layer, a hidden neurons layer and an output layer. CNN has achieved significant advances in recent years, with CNNs effectively increasing the flexibility and capacity of machine learning [63]. A CNN predictive model is made up of two parts, i.e., feature extraction and prediction. Compared to other models, the biggest strength of CNN lies in its convolution kernels, which focus only on local features and can fully extract the internal features of the data, boosting its accuracy. CNN have been applied to many SE research tasks. Liu et al. [63] exploited CNN to identify feature envy smells. The model they build had three layers with 128 kernels. The performance of their proposed

Table 5. The number of predictive models applied in per year.

Category	Family	Model Name	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	Total
Base learners	Regression	Logistic Regression	2	3	1	2	6	3	3	3	3	3	1	30
		Multiple linear regression			1				1					2
		Multinomial logistic regression										2		2
		Ordinary Least Squares							1				1	2
		Binary Logistic Regression							1					1
		Bayesian Logistic Regression							1					1
		Bayesian Ridge Regression											1	1
	Naive Bayes	Naive Bayes		3	3	4	4	3	4	3	2	3	1	30
		Bayesian Network		1	2			1	1		1	2	1	9
		Multinomial Naive Bayes										1		1
	SVM	SVM		2	3	3	5	2	2	2	1	3	4	27
		Linear SVM							1				1	2
		SMO							1	1				2
	Decision Tree	Decision Tree		2	1	2	1	1		1		1	4	13
		J48		1			1		4	2	1	1		10
		C4.5			1			1	2			1		5
		ADTree					1		1			1		3
		CART				1					3			2
		Alternating Decision Tree					1							1
		Recursive Partitioning								1				1
		K-means				1	2		1	1	1		1	7
		Expectation Maximization			1				1					2
		Spectral Clustering								1				1
Ensemble learning	Clustering technique	Fuzzy C- means								1				1
		Partitioning Around Medoids								1				1
		K-nearest Neighbor					1	1		1	1	1	1	6
		Nearest Neighbor Classifier				1	1		1	1	1			5
		nearest neighbor with non-nested generalization (NNGe)				1								1
	Discriminative Model	Markov Chains			1				1				1	3
		Conditional Random Field										2		2
	Rule-based	Conjunctive Rule						1	1					2
		Ripper							1	1				2
		Ridor							1					1
	Random Forest	Random Forest						1	3	6	4	6	3	23
		Tree Boosting										2	1	3
		AdaBoost						1				1		2
		XGBoost										1	1	2
		LogitBoost			1									1
	Bagging	gradient boosting											1	1
		Bagging				1	1		1	2		2		7
Neural network	Perceptron model	Multi-Layer Perceptron (MLP)					1		1		1	1	2	6
	MP Neuron model	Single-Layer Perception ANN							1		1		2	3
		AVNNet								1		1		2
		Siamese Neural Network (SNN)										1	1	2
	Deep learning techniques	LSTM										1	4	5
		RNN										1	1	2
		Bi-LSTM											1	1
		GRU											1	1
		Recurrent Highway Network										1		1
		CNN								1		2	1	4
		Deep Belief Network (DBN)								1				1

approach outperformed the state-of-the-art significantly in feature envy detection and recommendation. Xu et al. [149] present a deep-learning based approach to predict semantically linkable units in developers discussions in Stack Overflow. They formulated this problem as a multiclass classification problem, training a CNN to solve it. For classifying semantic relatedness, they used filters of five different window sizes, and each window size contained 128 filters to capture the most informative features.

Another approach is to use an LSTM model, a variant of the Recurrent Neural Network (RNN). This is specially designed for precessing sequential data. An LSTM model can capture the contextual information of sequences thanks to the recurrent nature of the RNN [160]. Computational units connected in each layer form a directed graph along a temporal sequence, which allows it to exhibit temporal dynamic behavior. Each LSTM unit contains an input gate, a memory cell, a forget gate, and an output gate. The gating mechanism of LSTM guarantees that the gradient of the long-term dependencies will not vanish [15].

Zhang et al. [160] utilized an attention-based Bi-LSTM model to detect anomalies in log events. Different from a standard LSTM, a Bi-LSTM can divide the hidden neuron layer into forward and backward, which allows it to capture more information of the log sequences. They confirmed the effectiveness of the proposed approach by comparing it with several traditional machine learning methods. Chen et al. [15] leveraged Bi-LSTM model as an emoji-powered learning approach for sentiment analysis. In their paper, the Bi-LSTM model contained two bi-directional LSTM layers and one attention layer, which treated the sentence vectors in Twitter and GitHub posts containing emojis as inputs and output the probabilities that instances contain each emoji.

5.2 Distribution of Different Predictive Models

In order to get a better understanding of predictive models, we summarize the frequency of models used in different software engineering tasks. Table 5 shows the number of cases for which different categories of predictive models and relevant studies in which they appear. Comparing with another two categories, the base learner is the most commonly used predictive models among all learners including eight different families, where there are 39 and 38 studies employing Logistic Regression and Naive Bayes to address software engineering problems in the last decade. SVM and Decision Tree are the next most frequently used predictive models with 31 and 35 studies using them, followed by Nearest Neighbor. Discriminative models and rule-based learners are also applied to software engineering tasks with high probability, including Markov Chains, CRF, Conjunctive Rule-based model, Ripper, and Ridor. It can be seen from Table 5 that the usage of base learners has remained stable between 2009 and 2019, indicating that these learners are still effective techniques in solving suitable problems, although they have been introduced some time ago.

There are 39 studies applying ensemble algorithms that can be classified into three categories, involving Boosting, Bagging, and Random Forest. 59% of the studies selected Random Forest as the predictive model. 9 studies leveraged Boosting techniques in their experiments, in which 3 used TreeBoosting and 2 studies each used AdaBoost and XGBoost. Another 2 studies adopted LogisticBoost and gradient boosting respectively to address specific problems. Bagging is another popular ensemble technique being used in 7 software engineering primary studies.

There are 29 studies that adopted Neural Network techniques, where 22 of the 29 studies used deep learning techniques, and most of these studies were conducted in 2018 and 2019. This suggests the application of deep learning to software engineering is exhibiting a booming trend in the last few years. In Table 5, the family of deep learning techniques, including LSTM, RNN, Bi-LSTM, GRU, and Recurrent Highway Network, is very popular with 10 studies applying them to software engineering problems. The second commonly used deep learning model is the CNN, with 4 studies using it. Deep Relief Network (DBN) was used only once among all selected studies. 7 studies adopted perceptron models, where Multi-Layer Perceptron was used in 6 studies. There are also 7 studies that employed MP neuron models, in which 3 used ANN, 2 used AVNNet, and another 2 studies used SNN.

5.3 Evaluation Metrics for Predictive Models

Table 6 lists the most widely used metrics in the primary studies, including a description, definition and the related studies. As Table 6 shows, Recall, Precision and F-measure are the most commonly used to evaluate the performance of predictive models, followed by Accuracy and AUC. 9 of the studies have used the probability of false alarm (pf) and 6 studies have used balance. ROC and MAE were applied in 5 studies each. Finally, only 4 studies selected G-measure and 3 selected specificity as evaluation measures.

Table 6. The definition and usage of common evaluation measures across the primary studies.

Measure	Description	Definition	#Studies
Recall	The proportion of positive cases that were correctly predicted	$\frac{TP}{TP+FN}$	96
Precision	The proportion of predicted positive instances that are correct.	$\frac{TP}{TP+FP}$	89
F-measure	Harmonic mean of precision and recall	$\frac{2 \times Precision \times Recall}{Precision + Recall}$	68
Accuracy	Proportion of correctly classified instances	$\frac{TP+TN}{TP+FP+TN+FN}$	32
Area Under the Curve (AUC)	The area under the receiver operating characteristics curve. Independent of the cutoff value		30
Probability of False Alarm (pf)	Proportion of negative instances incorrectly classified as positive	$\frac{FP}{FP+TN}$	9
Balance(bla)	The balance between pd (the probability of detection) and pf. The bigger the bla value is, the better the performance of learning model is.	$1 - \frac{\sqrt{(1-pd)^2 + (1-pf)^2}}{\sqrt{2}}$	6
ROC	A comprehensive indicators reflecting sensitivity and specificity of continuous variables		5
Mean Absolute Error (MAE)	The average of absolute errors, which can better reflect the actual situation of the predicted value error.	$\frac{1}{m} \sum_{i=1}^m (y_i - \hat{y})^2$	5
G-measure	Harmonic mean of pd and (1-pf)	$\frac{2 \times pd \times (1-pf)}{pd + (1-pf)}$	4
Specificity	The proportion of predicted negative instances that are correct	$\frac{TN}{FP+TN}$	3

Summary of answers to RQ2:

- (1) *Most of the primary studies applied base learners to tackle software engineering problems, where Logistic Regression and Naive Bayes are the most popular predictive models. Using deep learners in software engineering shows a thriving trend in recent years.*
- (2) *Recall, precision and F-measure are the three most commonly used evaluation metrics to evaluate the performance of different predictive models.*

6 RQ3: IN WHAT DOMAINS AND APPLICATIONS HAVE PREDICTIVE MODELS BEEN APPLIED?

Of the 145 studies in our survey, over 80% involved applying approaches using predictive models to a specific software engineering problems. In this section, we conduct an analysis on the distribution of different research domains to which predictive models were applied, and analyse the application of predictive models in each research topics.

To better understand the application distribution of predictive models, we categorise the software development life cycle (SDLC) into the following stages according to [9]: (1) Software Requirements; (2) Software Design and Implementation; (3) Software Testing and Debugging; (4) Software Maintenance; and (5) Software Engineering Professional Practice.

We referred to the definition of each category of the SDLC in [9] and conducted a comprehensive analysis of the problem addressed by each study and its practical application scenarios to determine which category the study belongs to.

6.1 Distribution in Different SDLC Domains

We categorize our selected primary studies based on research domains in Fig. 5 and illustrate the development trend of the main research domains in each year in Fig. 6. Table 7 lists the ranking of software engineering tasks that most studies have concentrated on.




Fig. 5. Research domains.




Fig. 6. Development trend of five main research domains.

As Fig. 5 shows, the primary research domain of most of the studies (50%) was software testing and debugging, followed by software maintenance (22%). Studies related to software engineering professional practice were 15%. Software design and implementation were the focus of 10% of the studies. Finally, only 8% of the papers focus on software requirements (3%). We notice that most of the primary studies worked on specific tasks with respect to software testing and maintenance. One potential reason is that these two research directions involve many predictive model-applicable tasks, such as defect prediction, code smell detection, and bug report management.

Table 7. Top Software Engineering Topics with a Minimum of Five Predictive Model Based Papers

Rank	Specific Task	Research Domain	#Studies
1.	Defect Prediction	Software Testing and Debugging	40
2.	Bug Report Management	Software Testing and Debugging	13
2.	Software Quality Management	Software Testing and Debugging	13
4.	Software Effort/Cost Prediction	Software Design and Implementation	10
5.	Developer Behavior Analysis	Software Engineering Professional Practice	9
6.	Code Clone Management	Software Maintenance	6
7.	Software Change Management	Software Maintenance	6

Fig.6 shows the development trend of the main research domains – Software requirements, Software development, Testing, Maintenance, quality assessment and Software engineering professional practice – over the last decade. Apart from 2009, there are a few studies related to software engineering professional practice in each year. Only 4 software requirement related studies using predictive models appeared in last decade. This indicates that software engineering professional practice shows a stable development trend, but there seems to be a lack of attention to predictive model research topics in software requirements. The proportion of studies in software maintenance shows an increase in recent years. We notice from Fig.5 that predictive models have not yet been widely used in

certain domains, such as software requirements and software design. This may perhaps motivate researchers to explore more scenarios in these SDLC domains suitable for application of predictive models.

Table 7 lists the ranking of specific topics where at least five predictive model related papers have appeared. We observe that five of the seven research topics belong to software testing and debugging (3) as well as software maintenance (2), which is consistent with the proportion of this research domain in Fig.5 and Fig.6. There are 40 publications employing predictive models in defect prediction, a research topic that most commonly to use predictive models as a solution, followed by but report management (13) and software quality management (13). The main contribution of 10 primary studies was to perform software effort / cost prediction by building predictive models. The number of publications working on code clone management is the same as that of software change management. Besides, 9 studies concentrated on developer behavior analysis (9).

6.2 Software Requirements

Requirements engineering (RE) is the process of defining, documenting, and maintaining requirements in software design process [9] and play a crucial role throughout systems' lifecycle. A requirement with high quality can notably improve development efficiency by avoiding reconstruction and rework.

6.2.1 Requirements classification. An effective classification of software requirements enables analysts to perform well-focused communication with developers and users as well as prioritize the requirement documents according to their importance [1].

Non-functional requirements (NFRs) classification: Anish et al. [3] conducted a study to identify, document, and organize Probing Questions (PQs) for five different areas of functionality into structured flows, called PQ-flow. They used Naive Bayes to identify Architecturally Significant Functional Requirements (ASFRs), used random k labelsets classifier (RAkEL) to categorize ASFR by types, and finally recommended PQ-flows.

Other requirements classification: Abad et al. [1] proposed a novel approach that extracts and classifies requirements-related knowledge to support analysts familiarity with different domains. Based on the intuition behind utilizing lexical association, they built a dynamic generative model to extract the relevant terms in documents and applied the rational kernels method as well as SVMs when classifying requirements.

6.2.2 Coordination of Requirements Detection. In large software projects, participants engage in parallel and interdependent tasks, resulting in work dependence and coordination needs. When developers remain unaware or do not obtain timely awareness of the coordination that is required to manage work dependencies, there is potential for software productivity and quality problems. Yang et al. [150] implemented a prototype tool for Nocuous Ambiguity Identification (NAI) in text. In their earlier work, they had focused on identifying the coordination ambiguity and anaphora ambiguity in requirement documents, and thus this study tried to detect another two types of ambiguity (i.e., nocuous and innocuous ambiguity). After extracting ambiguity instances and recognizing coordination constituents, they defined a set of heuristics (e.g., coordination matching, distribution similarity, and collocation frequency) based on word distribution and collocation and trained a binary classifier using LogitBoost algorithm to classify the input as nocuous or innocuous. Blincoc et al. [8] investigated what work dependencies a software development team needs to consider when establishing the coordination needs. They applied k-nearest neighbor machine learning algorithm to identify and supplement task properties that are indicative of the crucial coordination needs, contributing in the selection of the most critical coordination needs.

6.3 Software Design and Implementation

A software design is a description of the architecture of the software to be implemented, and algorithms or code management used. Below we introduce the application of predictive models in software design and for code development.

6.3.1 Software Architecture and Design. The architecture of a system describes the relationships between its main components and how they interact with each other. There are many factors involved in software architecture and design (e.g., quality, security, performance, etc.). It is a challenging task to balance different factors for designing a good architecture. To identify a set of quality concerns for maintaining security, reliability and performance of software systems, Mirakhorli et al. [80] proposed a cost-effective approach to automatically construct traceability links for architectural tactics. They applied a base learner as a tactic-classifier identifying all classes related to a given tactic. To reconstruct tactic-level traceability, these classes were mapped to their relevant tactic through two different classifiers, where one was trained using textual descriptions of each tactic and another using code snippets taken.

Architects construct software architectures by making a variety of design decisions that can satisfy quality concerns (e.g., reliability, performance, and security). However, the architectural quality decreased when developers modified code without fully understanding design decisions. In order to address this problem, Mirakhorli et al. [79] proposed a solution for detecting, tracing and visualizing architectural tactics in code. They utilized six machine learning algorithms (i.e., SVM, Decision Tree (C.45), Bayesian Logistic Regressions (BLR), AdaBoost, SLIPPER, Bagging) to train classifiers for detecting the presence of architectural tactics in source code. They monitored the architecturally significant code by mapping those relevant code segments into Tactic Traceability Patterns (tTPs) and notified developers when those segments were modified.

Gopalakrishnan et al. [31] present a bottom-up approach to identify architectural tactics. They adopted the Random Forest algorithm to capture relationships between latent topics and architectural tactics in source code of projects. Reman et al. [116] researched software architecture in another perspective. They performed a case study to investigate whether and how Hands-on software architects benefit the projects by using data analytics-based techniques in five large-scale industrial systems. To better understand what type of code architects write, they employed a classification algorithm [80] to identify functional comments and analyzed architects' code contributions.

6.3.2 Software Development Effort Estimation. In software development, Software effort estimation (SEE) is the process of predicting the most realistic amount of effort required to develop or maintain software [48]. SEE is an important activity in the software development process since it can provide beneficial suggestions for project plans, iteration plans, budgets, investment analyses, pricing processes and bidding rounds [48]. There are many studies related to effort estimation [16, 97]. To estimate SDLC, Nan et al. [97] formalized the nonlinear effect of management pressures on project performance as U-shaped relationships using Regression Models. They found that controlling for software process, size, complexity, and conformance quality, budget pressure, a less researched construct, has significant U-shaped relationships with development cycle time and development effort. Choetkietikul et al. [16] proposed a novel approach to provide automated support for project managers and other decision makers in predicting whether a subset of software tasks in a software project have a risk of being delayed. They used random forest as the local classifier on both local data and networked data.

More than half the literature on SEE has concentrated on comparisons of new estimation methods [17, 70, 76]. [20, 82] conducted comparative studies to investigate which estimation approach performs the best for software effort estimation. Whigham et al. [147] proposed an automatically transformed linear model (ATLM) as a suitable baseline model for comparison against software effort estimation methods. ATLM performed well over a range of different project types. Moreover, ATLM was able to be used with mixed numeric and categorical data and requires no parameter tuning. Dejaeger et al. [20] conducted a comparative study of data mining techniques for software effort estimation. Their techniques includes tree/rule-based models, linear models such as various types of linear regression, nonlinear models (MARS, multilayered perceptron neural networks, radial basis function networks, and least squares support vector machines), and estimation techniques that do not explicitly induce a model (e.g., a case-based reasoning approach). Their results showed that least squares regression in combination

with a logarithmic transformation performs best. Mittas et al. [82] proposed a statistical framework based on a multiple comparisons algorithm in order to rank several cost estimation models. They identified those which have significant differences in accuracy and clustered them into non-overlapping groups. They examined the predictive power of 11 models over 6 public domain datasets. Their results showed that very often a linear model is adequate enough, but not the optimal solution. Song et al. [129] aimed to present a synthetic data generator to alleviate the data scarcity problem of software effort estimation. They compared their approach with six state-of-the-art models to evaluate its performance. Maddila et al. [69] constructed a SEE model on an individual pull request level, helping developers trace their work items. Their methodology was deployed by several thousand developers across several software families. They observed from their deployment that 42.94% of developer time was saved by deploying their approach.

6.3.3 Software Process Evaluation. Software process evaluation is essential to improve software development and the quality of software products in an organization. Chen et al. [13] formulated this problem as a sequence classification task and introduced a novel semi-automated approach utilizing three well-known and representative classification algorithms (i.e., decision trees (C4.5), Naive Bayes (NB) algorithm, and Support Vector Machines (SVM)). Compared with traditional approaches, it was time-saving and avoided inaccurate evaluation results even if not each artifact can be accessed directly due to privacy or security concerns.

6.4 Software Testing and Debugging

Software testing is the process of verifying and validating whether a software or application is bug free, and its design and development meet technical and user requirements. Debugging is the process of identifying, analyzing and remove errors after the software fails to execute properly. Effective testing and debugging techniques can improve software quality with reduced effort consumption and improved performance [140].

6.4.1 Test Report Classification. Podgurki et al. [111] present an automated approach for classifying failure reports with closely related causes. Their proposed classification strategy involved supervised and unsupervised pattern classification, as well as multivariate visualization. Wang et al. [143] introduced a Local-based Active ClassiFication (LOAF) approach to classify crowdsourced test reports that reveal true faults. They adopted active learning to train a classification model with as few labeled inputs as possible in order to reduce the onerous burden of manual labeling. Compared with existing supervised machine learning methods, LOAF exhibited promising results on one of the largest Chinese crowdsourced testing platforms.

6.4.2 Test Case Quality Management. In software testing, test case quality is a critical factor. Good test cases can effectively improve software quality. On the contrary, constructing and running bad test cases may cost significant time and effort while still not finding defects or faults. There are many studies that aim to improve test case quality using predictive models.

Natella et al. [98] proposed a new approach to refine the “fault load” by removing faults that are not representative of residual software faults. They used a decision tree to classify whether a fault is representative or not based on software complexity metrics. Fard et al. [77] implemented a tool, called Testilizer, to automatically generating test cases for web applications. Since the problem of generating similar test suites was viewed as a classification task, they present a set of features and trained a classifier by leveraging human knowledge mining from human-written test suites.

Cotroneo et al. [18] proposed a method based on machine learning to adaptively combine testing techniques during the testing process. Their method contains an offline learning phase and an online learning phase. During offline learning, they first defined the features of a testing session potentially related to the techniques performance, and then used several machine learning approaches (i.e., Decision Trees, Bayesian Network, Naive Bayes, Logistic Regression) to predict the performance of a testing technique. During online learning, they adapt the selection

of test cases to the data observed as the testing proceeds. Yu et al. [154] proposed a technique that can be used to distinguish failing tests that executed a single fault from those that executed multiple faults. The technique combines information from a set of fault localization ranked lists, each produced for a certain failing test, and the distance between a failing test and the passing test that most resembles it.

6.4.3 Testing Applications. Malik et al. [72] applied three unsupervised approaches and one supervised approach to detect performance deviations. They first generated performance signatures which are minimal sets of performance counters, describing the essential characteristics of a System Under Test (SUT) for a given load test. They then classified each observed signature into passed or failed, and conducted a root-cause analysis of detected deviations. To reduce expensive calculation of mutation testing, Zhang et al. [158] proposed predictive mutation testing (PMT), that can predict the testing results without executing mutants. After identifying execution features, infection features and propagation features of tests and mutants, they selected Random Forest algorithm to construct a classification model for predicting whether a mutant can survive or be killed without mutant execution.

6.4.4 Software Defect Prediction. Software defect prediction techniques are employed to help prioritize software testing and debugging. These techniques can recommend software components that are likely to be defective to developers. Rahman et al. compared static bug finders and defect predictive models [114]. They found that in some cases, the performance of certain static bug-finders can be enhanced using information provided by statistical defect prediction models. Therefore, applying prediction models to defect prediction has become a popular research direction.

Lessmann et al. [58] proposed a framework for comparative software defect prediction experiments. They conducted a large-scale empirical comparison of 22 classifiers over 10 public domain data sets. Their results indicated that the importance of the particular classification algorithm may be less than previously assumed. This is because no significant performance differences could be detected among the top 17 classifiers. However, Ghotra et al. [28] doubted this conclusion and they pointed that the datasets Lessmann et al. used were both noisy and biased. Therefore, they replicated the prior study with initial its datasets as well as new datasets after cleansing. Their new experimental results demonstrated that some classification techniques were more suitable for building predictive models. They showed that Logistic Model Tree when combined with ensemble methods (i.e., bagging, random subspace, and rotation forest) achieves top-rank performance. Furthermore, clustering techniques (i.e., Expectation Maximization and K-means), rule-based techniques (Repeated Incremental Pruning to Produce Error Reduction and Ripple Down Rules), and support vector machine perform worse than other predictive models. Herzig et al. [36] conducted an analysis of the impact of tangled code changes in defect prediction. They used a multi-predictive model to identify tangled changes and they found that untangling tangled code changes can achieve significant accuracy improvements on defect prediction.

Many predictive model algorithms have several tunable parameters and their values may have a large impact on the model's prediction performance. Song et al. [130] proposed and evaluated a general framework for defect prediction that supports unbiased and comprehensive comparison between competing prediction systems. They first evaluated and chose a good learning scheme, consisting of a data preprocessor, an attribute selector and a learning algorithm. They then used the scheme to build a predictor. Their framework proposes three key elements for defect prediction models, i.e., datasets, features and algorithms. Tantithamthavorn et al. [136] conducted a case study on 18 datasets to investigate the performance of an automated parameter optimization technique, Caret, for defect prediction. They tried 26 classification techniques that require at least one parameter setting and concluded that automated parameter optimization techniques like Caret yield substantial benefits in terms of performance improvement and stability, while incurring a manageable additional computational cost.

Features, metrics or attributes are crucial for the building of a well-performed defect predictive model. A number of studies researched a lot about feature extraction for defect prediction [113, 126, 144]. Kim et al. [53] introduced a new technique for change-level defect prediction by using support vector machine. They are the first to classify

file changes as buggy or clean leveraging change information features. An obvious advantage of change-level classification is that predictions can be performed immediately upon the completion of changes. Shivaji et al. [126] investigated multiple feature selection techniques that are generally applicable to classification-based bug prediction methods. They found that binary features are better, and that between 3.12% and 25% of the total feature set yielded optimal classification results. Lee et al. [57] proposed 56 new micro interaction metrics (MIMs) by using developers' interaction information (e.g., file editing and selection events) in an Eclipse plug-in, Mylyn. To evaluate the effectiveness of MIMs, they built defect predictive models using traditional metrics, MIMS and their combinations. The evaluation results showed that MIMs allowed the performance of those models improve significantly. Yu et al. [153] initially focused on defect prediction for concurrent programs and proposed ConPredictor, a prototype tool that used four classification techniques to identify defects by applying a set of static and dynamic code metrics based on unique features of concurrent programs. Lin et al. [62] present an approach to identifying buggy videos from a set of gameplay videos. They applied three classification models – logistic regression, neural network and random forests – to determine the probability that a video showcases a bug. They observed that random forests achieved the best performance among classifiers. Afterwards, they identified key features in game videos and used random forest classifier to prioritize game videos according to the likelihood of each video containing bugs.

A good predictive model relies heavily on the dataset it learns from, which is also the case for models for software defect prediction. Many studies investigate the quality, such as bias and size, of datasets for defect prediction. Bird et al. [7] used predictive models to investigate historical data from several software projects (i.e., Eclipse and AspectJ), and found strong evidence of systematic bias. Rahman et al. [115] investigated the effect of size and bias of datasets on performance of defect prediction using logistic regression model. They investigated 12 open source projects and their results suggested that the type of bias has limited impact on prediction results, and the effect of bias is strongly confounded by size. Tantithamthavorn et al. [135] used random forest to investigate the impact of mislabelling on the performance and interpretation of defect models. They found that precision is rarely impacted by mislabelling while recall is impacted much by mislabelling. Moreover, the most influential features are generally robust to mislabelling.

For datasets of poor quality, researchers also have proposed several approaches to address their issues. Kim et al. [54] proposed an approach, named CLNI, to deal with the noise in defect prediction datasets. CLNI can effectively identify and eliminate noise. The noise eliminated from the training sets produced by CLNI was shown to improve the defect prediction model's performance. Nam et al. [95] proposed novel approaches CLA and CLAMI, which can work well for defect prediction on unlabeled datasets in an automated manner without any manual effort. Gong et al. [30] investigated the impact of class overlap and class imbalance problems on defect prediction, and then present an improved approach (IKMCCA) to solve those problems in order to improve defect prediction performance for within project defect prediction (WPDP) and cross project defect prediction (CPDP). Tantithamthavorn et al. [84, 134] conducted an empirical study to investigate how class rebalancing techniques impact the performance and interpretation of defect prediction. They also explored experimental settings that can help rebalancing techniques achieve the best performance in defect prediction. They concluded that using different metrics and classification algorithms allowed the performance of predictive models to vary, indicating that researchers should avoid class rebalancing techniques when deriving understandings and knowledge from models.

Datasets from many studies are unfortunately not always available due to privacy policies and other factors. To address the privacy problem, Peters et al. [105–107] measured the utility of privatized datasets empirically using Random Forests, Naive Bayes and Logistic Regression. Through this they showed the usefulness of their proposed privacy algorithm MORPH [105]. MORPH is a data mutator that moves the data a random distance, while not across the class boundaries. In a later work [106], they improved MORPH by proposing CLIFF+MORPH to enable effective defect prediction from shared data while preserving privacy. CLIFF is an instance pruner that deletes irrelevant examples. Recently, they again extended MORPH to propose LACE2 [107].

Deep learning based defect prediction: Recently, deep learning, as an advanced machine learning algorithm, has become widely discussed and applied including in Software Engineering. Several researchers have tried to improve the performance of defect prediction via deep learning [46, 144]. Jing et al. [46] proposed a cost-sensitive discriminative dictionary learning (CDDL) approach for software defect prediction. CDDL is based on sparse coding which can transform the initial features into more representative code. Their results showed that CDDL is superior to five representative methods, i.e., support vector machine, Compressed C4.5 decision tree, weighted Naive Bayes, coding based ensemble learning (CEL), and cost-sensitive boosting neural network. Wang et al. [144] leveraged Deep Belief Network (DBN) to automatically learn semantic features from token vectors extracted from programs Abstract Syntax Trees. Their evaluation on ten open source projects showed that learned semantic features significantly improve both within-project defect prediction and cross-project defect prediction compared to traditional features.

Cross project defect prediction: Cross-project defect prediction is a topic of growing research interest. It uses data from one project to build the predictive model and predicts defects in another project based on the trained model so that it can solve the problem that there is no sufficient amount of data available to train within a project, such as a new project. Some studies concentrated on cross-project defect predictive models on a large scale [167].

To improve the performance of cross-project defect prediction, researches have tried several techniques [44, 94, 133]. Nam et al. [96] proposed a novel transfer defect learning approach, TCA+, extending a transfer learning approach Transfer Component Analysis (TCA). TCA+ can provide decision rules to select suitable normalization options for TCA of a given source-target project pair. In a later work, they addressed the limitation that cross-project defect prediction cannot be conducted across projects with heterogeneous metric sets by proposing a heterogeneous defect prediction approach. Jiang et al. [44] proposed an approach CCA+ for heterogeneous cross-company defect prediction. CCA+ combines unified metric representation and canonical correlation analysis and can achieve the best prediction results with the nearest neighbor classifier.

Zhang et al. [156] found that connectivity-based unsupervised classifiers (via spectral clustering) offer a viable solution for cross-project defect prediction. Their spectral classifier ranks as one of the top classifiers among five widely-used supervised classifiers and five unsupervised classifiers in cross-project defect prediction. Due to the impact of heterogeneous metric sets, current defect prediction techniques are difficult to use in CPDP. To address this problem, Nam and Kim [94] trained a predictive model for heterogeneous defect prediction (HDP) with heterogeneous metric sets by applying metric selection and metric matching techniques. Zhou et al. [165] noticed that most CPDP models were not compared against simple defect predictive models, and thus they aimed to investigate whether CPDP models really performed well compared against simple models, ManualDown and ManualUp. To their surprise, they found that the performance of ManualDown and ManualUp were superior to most of existing CPDP models.

Just-in-time (JIT) defect prediction: Compared with traditional defect predictions at class or file level, Just-in-Time (JIT) defect prediction is of more practical value for participants, which aims to identify defect-inducing changes. Many studies focused on JIT defect prediction by employing SZZ approach [24, 49, 50, 126]. In order to identify bug-introducing changes, SZZ first detects the bug-fixing changes whose the change log contains bug identifier. Among the bug-fixing changes, SZZ identifies the modified the buggy lines so that the bug can be removed. SZZ then traces the code change history to search for potential bug-introducing changes that may introduce buggy lines. Finally, the remaining changes are deemed as bug-introducing after SZZ eliminates the incorrect ones from the initial set of potential bug-introducing changes. Yang et al. [151] used the most common change metrics from various source to build unsupervised and supervised models to predict defect-inducing changes. Comparing these models under cross-validation and time-wise-cross-validation, their results showed that the simple unsupervised models performed better than state-of-the-art supervised models for JIT defect prediction.

Failure prediction: Failure prediction is critical to predictive maintenance since it has the ability to prevent maintenance costs and failure / bug occurrences. Salfner et al. [119] conducted a survey on online failure prediction

approaches. Online failure prediction is based on runtime monitoring and a variety of models and methods that use the current state of a system and, frequently, the past experience as well. This study classified the existing online failure prediction techniques into three main categories: failure tracking-based, error reporting-based, and symptom monitoring-based approaches. Yilmaz and Porter [152] classify measured executions into successful and failed executions in order to apply the resulting models to systems with an unknown failure status. They further applied their technique to online environments. Ozcelik and Yilmaz [100] proposed a online failure prediction approach for combining hardware and software instrumentation. In the training phase, they applied global and frequency filtering to pick up candidate functions from historical data as the input of the proposed approach and created a predictive model for identifying the best performing functions that best distinguish failing executions from passing executions. These best performing functions were referred to as seer functions. Then their proposed model makes a binary prediction for each seer function into passing or failing during monitoring phase. To address the understudied problem of inferring concurrency-related documentations, Habib et al. [33] developed a novel tool, TSFinder, to automatically classify classes as thread-safe or thread-unsafe with a combination of static analysis and graph-based classification. Yilmaz and Porter [152] present a hybrid instrumentation approach to distinguish failed executions from successful executions with consideration of potential cost-benefit tradeoffs. After collecting program spectra and augmenting the underlying data with low cost, they trained a classification model with J48 algorithm to classify program executions.

6.4.5 Bug Report Management. Bug reports are essential for any software development, which typically contain a detailed description of the failure and occasionally hint at the location of the fault in the source code in the form of patches or stack traces. They also allow users to inform developers of problems encountered while applying software applications.

Duplicate bug report detection: To reduce redundant effort, there has been much research into identification of duplicate bug reports [131, 132]. Sun et al. [132] used discriminative models to identify duplicates more accurately. They used a support vector machine (SVM) with linear kernel based on 54 text features. In a later work [131], they more fully utilized the information available in bug reports including not only textual content and description fields, but also non-textual fields (e.g., product, component, version) and proposed a retrieval function REP to identify duplicated reports. A two-round stochastic gradient descent was applied to automatically optimize REP in a supervised learning manner.

Bug report quality management: Zanetti et al. [155] proposed an efficient and practical method to identify valid bug reports i.e., the bug reports that refer to an actual software bug, are not duplicates, and contain enough information to be processed right away. They used support vector machine to identify valid bug reports based on nine network measures using a comprehensive data set of more than 700,000 bug reports obtained from the BUGZILLA installation of four major OSS communities. To effectively analyze anomaly bug reports, Lucia et al. [67] proposed an approach to automatically identify and refine bug reports by the incorporation of user feedback. They first listed the top few anomaly reports from the list of reports generated by a tool in its default ordering. Based on feedback of users who either accepted or rejected each of the reports, their tool adopted nearest neighbor with non-nested generalization (NNGe) to automatically and iteratively refine a classification model for anomalies and resorted the rest of the reports.

Bug report assignment and categorization: Considering the impact of bug report assignment on the efficiency of development process, making a proper and effective assignment of each bug report is a crucial task. This can save much time and effort on bug-fixing activity. Jeong et al. [42] introduced a graph model based on Markov chains, which captures bug tossing history. They showed that the accuracy of bug assignment prediction is improved using naive bayes with tossing graph. Anvik et al. [4] presented a machine learning approach to create recommenders that assist with a variety of decisions aimed at streamlining the development process. They built three different development-oriented recommenders by applying six predictive models to suggest which developers might fix

errors in the bug reports, which product components a report might pertain to, and which developers on the project might be interested in following the report. Jonsson et al. [47] conducted a study to evaluate machine learning classification based automated bug assignment techniques. They combined an ensemble learner Stacked Generalization (SG) with several classifiers and evaluated their performance with over 50,000 bug reports from two companies. Peters et al. [108] present a framework to filter and remove non-security bug reports containing security related keywords by training five predictive models – Random Forest, Naive Bayes, Logistic Regression, Multilayer Perceptron, and K Nearest Neighbor – on Four Apache projects and Chromium involving 45,940 bug reports. They found that their framework mitigates the class imbalance issue and reduces the number of mislabelled security bug reports by 38%. Proper bug report assignment also relies on proper bug report categorization. There are a number of studies that propose techniques to categorize bug reports. Among them a popular research area is reopened bug report prediction. Zimmermann et al. [168] characterized and predicted which bugs get reopened. They first qualitatively identified causes for bug reopens, and then built logistic regression model to predict the probability that a bug will be reopened.

Bug report-based bug fixing tasks: A number of studies leverage bug report contents to address other bug fixing related tasks. Guo et al. [32] performed an empirical study to characterize factors that affect which bugs get fixed in Windows Vista and Windows 7. They concentrated on investigating factors related to bug report edits and relationships between people involved in fixing bugs, and built a statistical model using logistic regression to predict the probability that a new bug will be fixed. Kim et al. [52] proposed a two phase predictive model to identify the files likely to be fixed by analyzing bug reports contents. In order to collect predictable bug reports, they checked whether the given bug report contains sufficient information for prediction (predictable or deficient) in the first phase. Zhang et al. [157] proposed a Markov-based method for predicting the number of bugs that will be fixed in future. For a given bug report, they also constructed a KNN classification model to predict bug-fixing effort based on the intuition of similar bugs requiring similar time cost.

6.4.6 Software Quality Assessment. With the continuous release of new software versions, developers have been committed to maintaining a high level of software quality during the whole of SDLC. Generally, Software quality can be defined as “the degree to which a software product meets established requirements” and it involves many aspects, such as desirable characteristics of software products and process, techniques, or tools used to achieve those characteristics. In this section, we introduce the studies using predictive models for assessing software quality.

Vulnerability Detection: An indicator being commonly used for evaluating software quality is software vulnerability. There are many studies on vulnerability detection using predictive models [19, 120, 123, 125, 159]. Shin et al. [125] conducted an empirical case study to investigate the impact of software metrics obtained in the early SDLC on predicting vulnerable code. They collected complexity, code churn, and developer activity metrics from two large-scale projects (Mozilla Firefox and the Linux kernel) and trained five machine learning models to identify vulnerabilities by applying these metrics. The experimental results showed that Naive Bayes achieve the best performance among other classifiers. Scandariato et al. [120] present an approach to predict which components of a software application contain security vulnerabilities by training five machine learning models. Based on [123], Zhang et al. [159] utilized deep learning techniques to detecting SQL injection vulnerabilities in PHP code. They observed that a classifier trained using CNN outperformed another one trained by Multilayer Perceptron (MLP). Dam et al [19] built a Long Short Term Memory deep learning-based model to automatically identify key features that predict vulnerable software components. Their experiments on Firefox and 18 Android apps showed that LSTM performs better than other predictive models for both within-project and cross-project vulnerability detection.

Malware detection: Some software, especially mobile apps, can include significant malware. Generally, malware includes viruses, worms, trojans, key loggers and spyware, and are harmful at diverse severity scales.

Malware can lead to damages of varying severity, ranging from spurious app crashes to financial losses with malware sending premium-rate SMS, and to private data leaks. A number of studies have aimed to better detect malware by using predictive models [5, 12].

Chandramohan et al. [12] proposed and evaluated a bounded feature space behavior modeling (BOFM) framework for scalable malware detection. They first extracted a feature vector that is bounded by an upper limit N using BOFM, and then trained an SVM to detect malware. Avdiienko et al. [5] compared data flow of 2,866 benign applications against those in malicious applications to capture differences. They noticed that malicious applications treat sensitive data differently than benign applications, which can be used to identify malicious applications. Therefore, they present MUDFLOW prototype to flag malware by using abnormal data flow. Since there is no studies to detect malicious and aggressive push notification advertisement, Liu et al. [65] provided a taxonomy of push notifications and present a crowdsourcing-based approach to automatically detect aggressive push notifications in Android apps. Their approach used a guided testing approach to record push notifications and flagged the malicious ones by analyzing runtime information.

Product Quality Assessment: High quality software products can improve user experience and reduce developers' effort in maintenance. Quality assessment is an important task to discover and fix new bugs that appear in products timely. Miller et al. [87] investigated whether biometrics are able to determine code quality concerns. They conducted a field study with ten professional developers and further adopted machine learning classifiers to predict developers' perceived difficulty of code elements while working on a change task. Mills et al. [78] proposed a solution to the problem of predicting query quality by training a machine learning classifier based on 28 query measures. They applied a Random Forest algorithm to 12 open source systems implemented in Java and C++, defining seven post-retrieval query quality properties, which extended 21 pre-retrieval properties proposed by previous work [34, 35]. In order to improve the efficiency of developers in handling test reports, Chen et al. [14] present a framework to assess the quality of crowdsourced test reports. They trained a classifier using logistic regression to predict the quality of reports. The proposed framework achieved good performance in terms of precision, recall and F-measure.

6.5 Software Maintenance

Software maintenance is an integral stage of the software life cycle. No software product is ever completely finished and all need some form of ongoing maintenance. We introduce representative applications of predictive models for software maintenance.

6.5.1 Configuration Management. When developers or researchers build predictive models in different application scenarios, the effect of model configuration, i.e., the parameters, specify their behavior. There are several studies focusing on the effect of parameters in predictive models or which parameter values lead to the best performance [127, 128, 138].

Thomas et al. [138] investigated the effectiveness of a large space of classifier configurations, 3,172 in total, and present a framework for combining the results of multiple classifier configurations since classifier combination has shown promise in certain domains. Due to several limitations of the combinatorial interaction testing (CIT) approach, Song et al. [127] implemented an iterative learning algorithm called iTree, which effectively searched for a small part of configurations that closely approximated the effective configuration space of a system. Based on their previous work [128], the key improvements of are based on the use of composite proto-interactions, i.e. a construct that improves iTrees ability to correctly learn key configuration option combinations. This in turn significantly improves iTrees running time, without sacrificing effectiveness. Nair et al. [92] present a rank-based approach to rank software configurations and identify the optimal value without requiring exact performance values. Experiments were conducted on 21 scenarios for evaluating effectiveness of their strategy, and results showed

that rank-based approach allows building an accurate performance model with very few data instances, helping developers to significantly reduce the cost of building models.

6.5.2 Software Change Management. Change management involves tracking and managing changes to artifacts, such as requirements, change requests, bug reports, source code files, and other digital assets. It's critical for effective application development. With the rapid evolution and change of software, effective change management has become a major challenge in the industry, and thus many studies have worked on identifying changes, analyzing changes, tracing changes, assessing security or metrics using predictive models [23].

Padhye et al. [101] present a prototype tool for highlighting code changes in order to improve development efficiency. They explored the use of various techniques including Naive approach (TOUCH-based subscription strategy) and machine learning algorithms (Conjunctive Rule, J48 and Naive Bayes) to model code relevance. Their proposed tool can highlight changes that developers may need to review in order to personalize a developers' change notification feed.

Liu et al. [66] proposed a novel learning-based approach that can predict software licenses as software changes by applying linear-chain conditional random field (CRF) learning algorithm. To demonstrate its effectiveness, they evaluated their approach on 700 open source projects. Misirli et al. [81] analyzed high impact fix-inducing changes (HIFCs). They used a Random Forest-based algorithm to identify HIFCs and determine the best indicators of HIFCs on six large open source projects. They also present a measure to evaluate the impact of fix-inducing changes. Falessi et al. [23] introduced new Requirements to the Requirements Set (R2RS) family of metrics based on some intuitions. In order to evaluate the usefulness of the proposed metrics, they applied five classifiers to predict the set of classes impacted by a requirement over 700,000 classes. Horton and Parin et al. [37] proposed V2, a strategy to detect discrete instances of configuration drift in order to determine whether a code snippet is out-of-date and where code snippets that use APIs experience a breaking change.

6.5.3 Code Clone Management. Code clones have always been a double-edged sword in software development. On one hand, developers reuse the existing code snippets to complete development tasks, largely reducing coding effort. On the other hand, creating clones in source code may introduce new defects or bugs, which can lead to extra maintenance effort to ensure consistency among cloned code snippets. Since software clones have a clear influence on software maintenance and evolution, many studies [85, 90, 118, 139, 145, 146] proposed various approaches to detect and management code clones in source code.

Code Clone Detection: Saini et al. [118] implemented a novel code clone detection approach referred as to Oreo, which can be capable of detecting harder-to-detect clones (i.e., semantic clones) in the Twilight Zone. They employed five different models to detect clones and selected siamese architecture neural network as the final choice. They evaluated the recall of Oreo on vBigCloneBench, and Oreo achieved both high recall and precision. Nafi et al. [90] proposed a framework which can detect cross language clones (CLCs) with no need to generate the intermediate representation of source code. Their framework analyzed the different syntactic features of source code across different programming languages and applied siamese architecture neural network to detect CLCs.

Code Clone Analysis: Prior studies had demonstrated that neither the percentage of harmless code clones nor that of harmful code clones is negligible [29]. In order to assist developers in avoiding the negative impacts of harmful code clones and taking advantage of the benefits of clones, they present a novel approach that can automatically predict the harmfulness of an intended cloning operation after performing copy-and-paste by using Bayesian Networks. Since only a small part of code clones experience consistent changes during software evolution history, Wang et al. [146] defined a code cloning operation as consistency-maintenance-required if a code clone experience consistent changes in their evolution history, and leveraged Bayesian Networks to automatically predict whether a code cloning operation needs consistency maintenance when developers performing copy-and-paste operations. They evaluated the effectiveness of their approach on four projects under two usage scenarios. In order to perform the efficiency of clone management efforts, Thongtanunam et al. [139] conducted an empirical study on

six open-source Java systems to investigate the life expectancy of clones. They found that 30% to 87% of clones were short-lived, and these short-lived clones were changed more frequently than long-lived clones throughout their lifetime. They also applied Random Forest classifier to predict the life expectancy of newly-introduced clones. Mostaeen et al. [85] present a tool called CloneCognition that can automate the laborious manual validation process by using Artificial Neural Network (ANN). Their tool showed promising performance when compared with state-of-the-art clone validation techniques.

6.5.4 Code Smell Detection. Code smell refers to the symptoms of poor design and implementation choices in source code, which possibly indicate deeper problems for further development, maintenance, and evolution of software. A large number of studies have been proposed to automatically detect various types of code smells since it is tedious and time consuming to manually identify code smells [75].

Fontana et al. [26] applied 16 different machine learning techniques on large-scale code smell instances. By comparing the performance of various machine learning algorithms, they found that all algorithms performed well on cross-validation datasets. Their experimental results showed that the highest performance was achieved by J48 and Random Forest, while SVMs had the worst performance in code smell detection. Palomba et al. [102] performed an empirical study to investigate the relationship between community smells and code smells. They found that community smells result in the increasing intensity of code smells in projects. Liu et al. [63] proposed a neural network-based classifier that detects feature envy without the need for manually selecting features. Evaluation results showed that the proposed approach significantly improved the state-of-the-art in identifying feature envy smells.

6.5.5 API Issue Classification. In order to identify API issue-related sentences in Stack Overflow (SO) Posts and to classify SO posts concerning API issues, Lin et al. [60] introduced a machine learning based approach called POME. POME classifies sentences related to APIs in Stack Overflow, using natural language parsing and pattern-matching and determined their polarity (positive vs negative). Evaluation results showed that POME exhibited a higher precision compared with a state-of-the-art approaches. Similarly, Ahasanuzzaman et al. [2] developed a supervised learning approach with a Conditional Random Field (CRF). They also performed an investigation to identify important features and test the performance of the proposed technique for classifying issue sentences. Petrosyan et al. [110] classified fragmented tutorial sections that help developers to understand a subset of API types. They applied a supervised text classification based on structural and linguistic features. applying supervised text classification according to structural and linguistic features. Nam et al. [93] conducted studies on boilerplate code and investigated what properties make code be considered as “boilerplate”. They then proposed a novel approach automatically mine boilerplate code candidates from API client code repositories.

6.5.6 Log Analysis. Logging provides visibility into the health and performance of an application and infrastructure stack. Some studies concentrated on log analysis in order to enable development teams and system administrators to more easily diagnose and rectify issues. Russo et al. [117] proposed an approach for classification and prediction of defective log sequences. They used three well-known SVMs – radial basis function, multilayer perceptron and linear kernels – to predict and fit log sequences containing defects with high probability. The approach achieves comparable accuracy as other log analysis techniques. Li et al. [59] studied reasons of log changes and proposed an approach to determine whether log change suggestions are required while committing a code change. Zhang et al. [160] present LogRobust, a log-based anomaly detection approach, using an attention-based Bi-LSTM model. They evaluated LogRobust on Hadoop and the results demonstrated that LogRobust can identify and handle stable log events and sequences with high accuracy.

6.5.7 Code Comment Analysis. Comments in source code explain the purpose of the code and help with code understanding during development and maintenance. Huang et al. [40] present a classification approach to detect self-admitted technical debt (SATD) in code comments. They applied feature selection to select valuable

features and trained a composite classifier by combining multiple classifiers for identifying SATD comments in open source projects. For better understanding the goal and target audience of code comments, Pascarella et al. [103] investigated how diverse Java software projects used code comments, and applied machine learning algorithms to automatically classify code comments at line level.

6.5.8 Software Repository Mining. Software repositories, such as GitHub and Q&A sites, contain a large amount of knowledge related to software development. Mining these repositories is able to recover valuable knowledge that can help developers address difficulties when programming. However, the quality of a high proportion of questions and answers in repositories is still a concern. To ensure that researchers reach realistic and accurate conclusions, Munaiah et al. [88] trained a Random Forest classifier for sieving out the noise in GitHub repositories. They evaluated their framework on 200 repositories with known ground truth classification. Xu et al. [149] applied deep learning techniques to predict Semantically Linkable Knowledge units in Stack Overflow. They formulated the problem of identifying knowledge units as a multi-class classification task and trained a convolutional neural network (CNN) model by exploiting informative word-level and document-level features. Prana et al. [112] manually annotated 4,226 README files from GitHub repositories and designed a multi-label classifier to classify these files into eight different categories.

6.5.9 App Review Mining. In order to capture the meaning of a sentence at a higher level of abstraction, Jha and Mahmoud [43] applied frame semantics to generate low-dimensional representations of text and used SVM and Naive Bayes to classify app store reviews into three categories of actionable software maintenance requests (i.e., bug reports, feature requests, and otherwise), enhancing the predictive capabilities and reducing the chances of overfitting. Martens et al. [73] performed a survey of 43 fake review providers to study the significant differences between fake reviews and non-fake reviews. They then implemented seven classifiers to automatically detect fake reviews in app stores. The experimental results showed that the Random Forest algorithm achieved the best performance compared to other models.

6.6 Software Engineering Professional Practice

Software engineering professional practice is related to the skills, knowledge and attitudes that software engineers need to possess to practice software engineering in a professional, responsible and ethical manner. The study of professional practices includes the areas of technical communication, group dynamics and human aspects including psychology, emotions, team dynamics, social and professional responsibilities [122]. On the other hand, the goal of software development is to meet user desires or needs [9]. We introduce some key related studies into these developers' and users' perspectives that use predictive models.

6.6.1 Developer Behavior Analysis. Meneely et al. [74] performed an empirical and longitudinal case study of a large Cisco networking product over a five year period. They examined statistical correlations between monthly team-level metrics and monthly product-level metrics. Their linear regression predictive model based on team metrics was able to predict the products post-release failure rate within a 95% prediction interval for 38 out of 40 months. Canfora et al. [11] proposed an approach named Yoda, that aimed to identify and recommend mentors in software projects by mining data from mailing lists and versioning systems. Fritz et al. [27] implemented a novel approach to detect when software developers are experiencing difficulty while they work on their programming tasks, and stop developers' behavior before they introduce bugs into the code. They classify the difficulty of code comprehension tasks using data from psycho-physiological sensors and chose Naive Bayes as the classification algorithm because its training can be updated on-the-fly.

Muller et al. [86] investigated developers emotions, progress and applied biometric measures to classify them in the context of software change tasks. They trained a J48 decision tree to distinguish between positive and negative emotions based on biometric measurements. Bacchelli et al. [6] presented an approach to classify email content at

line level. Their technique fused naive Bayes algorithm with island parsing to perform automatic classification of the content of development emails into five language categories: natural language text, source code fragments, stack traces, code patches, and junk. Their technique can help developers subsequently apply ad hoc analysis techniques for each category. Later, Sorbo et al. [21] proposed a semi-supervised approach named DECA (Development Emails Content Analyzer) to mine intention from developer emails. DECA uses Natural Language Parsing to classify the content of development emails according to their different purposes (e.g., feature request, opinion asking, problem discovery, solution proposal, or information giving), identifying email elements that can be used for specific tasks. They showed the superiority of DECA to traditional machine learning techniques.

Zhou et al. [162] measured, understood, and predicted how the newcomers involvement and environment in issue tracking systems affects their odds of becoming a long term contributor. They constructed nine measures of involvement and environment based on events recorded in an issue tracking system and used logistic regression model to predict long term contributors. Wood et al. [148] conducted an empirical study with 30 professional programmers and trained a supervised learning algorithm to identify speech act types in developers' conversations in order to obtain useful information for bug repair.

6.6.2 User Behavior Analysis. Murukannaiah et al. [89] believed that an application cannot benefit users by employing location information, but rather a position, the abstraction of location, such as restaurant, park, home, and office. They referred this level of location abstraction to as "place". Therefore, they proposed the Platys framework as a way to address the special challenges of place-aware application development. They collected place labels and Android phone sensor readings from 10 users, and applied Platys to learn each users places by combining active learning and semi-supervised learning. Zheng et al. [161] implemented a tool referred to as iFeedback to perform real-time issue detection by analyzing user feedback texts. iFeedback trained a sophisticated 2-class classification model by employing XGBoost, which allowed it to detect fast system anomalies quickly and accurately. Sharma et al. [124] employed a two-stage classification approach to identify specialized software gurus by analyzing tweets. Their approach can be applied to identify specialized gurus in four domains (i.e., Android, JavaScript, Linux and Python) on a dataset of 86, 824 Twitter users for evaluating its effectiveness.

Summary of answers to RQ3:

- (1) *The application of predictive models in 22 research topics throughout the SDLC were summarized.*
- (2) *Software testing and debugging is the domain in which 50% of predictive models have been applied in the selected primary studies. 40 studies use predictive models for software defect prediction.*
- (3) *Only 4 of the studies worked on requirements classification using predictive models. Researchers may want to focus more attention on using predictive models in analyzing software requirements and testing.*

7 THREATS TO VALIDITY

Publication Bias. Publication bias is the issue of publishing more positive results over negative ones. Claims to reject or support a hypothesis will be biased if the original publication is suffering from bias. A tendency toward positive study outcomes rather than others leads to biased and even possibly incorrect conclusions, while some preferences in publishing are useful. Studies with a null result might not always be worse than studies with significant positive results, although significant results have a statistically higher chance of getting published [55]. In this paper, we select six top and prevalent software engineering venues. Thus, to some extent, the publications included in these venues are high-quality, which can reduce or eliminate the influence of publication bias on conclusions.

Search Terms. Finding all relevant primary studies is still a challenge to any survey or literature review-based study. To tackle this issue, a detailed search strategy was prepared and performed in our research. Search terms were constructed with different strings identified by checking titles and keywords from relevant publications already known to the authors. Alternative synonyms and spellings for search terms were then modified by consulting an expert. These procedures provided a high confidence that the majority of the key studies were identified.

Study Selection Bias. The publication selection process was carried out in two stages. In the first stage, studies were excluded based on the title and abstract independently by two researchers with extensive experience in software engineering. We conducted a pilot study of publication selection process to place a foundation in order to better understand the inclusion/ exclusion criteria, and evaluated inter-rater reliability to mitigate the threat emerged from the researchers personal subjective judgment. When two researchers could not reach an agreement on a study, a third researcher was consulted. In the second stage, studies were eliminated based on the full paper. Thus there is little possibility to miss relevant studies through this well-established study selection process.

Data Extraction. Data extraction was conducted by two of the researchers, and the data extracted from the relevant studies was rechecked by the other researchers. Disagreements in the data extraction process were discussed and solved after the pilot data extraction so that the researchers could complete the data extraction process following the refinement of the criteria, improving the validity of our analysis.

8 CHALLENGES AND OPPORTUNITIES

A number of research challenges and opportunities relating to the use of predictive models in software engineering remain, requiring further investigation. In this section, we discuss the limitations of using predictive models in different SE research areas and conclude four key findings pertaining to challenges when building these models. We also present several recommendations for future work.

8.1 Challenges

8.1.1 Challenge 1: Quality of Datasets. Source datasets for training and testing are critical for predictive model-based approaches. The most significant challenge for most studies is whether the information in the datasets available reflect ground truth about the related software engineering tasks. With different research areas, there are a large number of datasets that can be used for experiments. However, the datasets vary much in quality, including bias, noise, size, imbalance and mislabelling. Any of these will influence the effectiveness of the predictive models used. For example, a dataset with much noise may cause underfitting in building predictive models. A dataset with small size may cause overfitting in building predictive models. A heavily imbalanced dataset may even fail to build useful models at all.

Data imbalance: The datasets are usually not distributed evenly [45]. In some cases, over 90% of the total data have data imbalance problem. Some datasets even contain a large proportion of false positive instances. This issue causes the class imbalance problem, which has a large influence on the performance of predictive models. Learning a minority class is difficult when a classifier is trained on imbalanced data. Therefore, the classifier is skewed toward the majority class resulting in a lower rate of detection. Several studies proposed different methods to address this problem, such as instance re-weighting [91], data re-sampling [83], and selective learning strategy [39].

Noise in datasets: Compared with the acceptable range of the feature spaces, data instances having excessive deviation can be referred to as outliers [38]. Classifiers trained on noisy data with many outliers (e.g., incorrect, missing information) are expected to be less accurate. Since the problem of noisy data is inevitable from time to time when dealing with certain datasets such as software repositories, the usage of such noisy data is likely to pose a threat to the validity of many studies [54]. Prior studies demonstrated that noisy may creep into datasets and impact predictive models when collecting data carelessly [54]. There are two lines of related studies focusing on

this problem. The first conducted a comprehensive investigation of the impact of noisy data on models performance [56]. The second line of studies proposed different textual similarity approaches to address this challenge [56].

Size of datasets: Most of selected studies mentioned the problem of generality of their proposed approaches as an external threat due to lack of enough data. Using representative datasets is one of crucial factors towards performing experiments [115]. For instance, some studies trained their predictive models using source code in a single programming language such as C, Java or Python. This means the training classes cannot comprehensively cover all possible patterns and may prevent those algorithms from addressing corresponding tasks in all cases. In addition, overfitting may affect models' ability when training models without large-scale datasets. A number of studies have worked on alleviating this problem by constructing large-scale datasets from various software repositories [64]. However, a sufficiently large scale volume of data in some communities is not yet available to be used.

Data Privacy: Data privacy issues arise when the confidentiality of the data is a concern for the project owners. This issue in turn may cause the owners to not contribute to the pool of available data even though such data might contribute to further research efforts. Privacy is considered by Peters et al. [106] in the context of Cross-Company Defect Prediction (CCDP). They obfuscate the data in order to hide the project details when they are shared among multiple parties.

Data heterogeneity: The outcome of predictive models can be easily influenced by the similarity of source and target data distributions, and this problem is called data heterogeneity. As expressed by Canfora et al. [?], most software projects are heterogeneous and have various metric distributions. Moreover, certain context factors (e.g., size, domain, and programming language) have a huge impact on data heterogeneity. Machine learning techniques can work well under the assumption that source and target data have the same distribution. Some studies [141, 167] present different approaches such as filtering (F), transformation (DT), normalization (N) and feature matching to tackle this problem.

8.1.2 Challenge 2: Feature selection. As many candidate features are introduced to better represent various research domains, many studies extracted and used a greater variety of features to train predictive models, increasing the size of research space of features. However, not all features used are beneficial to improving a predictive model's performance [19]. First, some interdependent features do not fit when applied to build models. Another reason is that using too many features may result in overfitting. Therefore, how to select the most suitable features in a huge research space has become a critical challenge. Feature selection techniques ensure that useless features are removed. Some studies [113, 126] employed such techniques in order to optimize performance and scalability. A good feature extractor should meet two main principles. First of all, it should be automatic so that it won't cost too much manual effort. Second, it can reduce feature dimension while keeping feature quality. That is, given a specific task, the feature extractor can identify and preserve relevant features and remove irrelevant features to improve the performance of predictive models. Automatic feature selection may even be possible in some research directions [19].

8.1.3 Challenge 3: Evaluation metric selection. Diverse evaluation metrics have been used to evaluate the effectiveness of different models. However, evaluation metrics may introduce bias. For example, [40, 165] highlight the problem of the suitability of the chosen evaluation metrics. Most primary studies solve this challenge by utilizing standard metrics. Instead of using standard metrics such as precision, recall and F-measure, many studies [27, 67, 91, 162] have defined new evaluation measures to evaluate the performance of their proposed frameworks when addressing the same problem. This makes it difficult to evaluate the performance of their frameworks. On the other hand, the measures used for evaluation models keep changing thanks to new techniques being found. Therefore, it is a challenging task to achieve optimal selection of evaluation measures.

8.1.4 Challenge 4: Guidelines for the selection of suitable predictive models. From our survey, we see that there are various software engineering tasks leveraging predictive models. However, different predictive models fit better for different tasks. With various tasks and various predictive models, there is a need to have reliable guidelines for how to select a good predictive model for a specific tasks. Among the reviewed papers, some studies have proposed frameworks for specific tasks to make comparison of various predictive models [20, 28, 82, 84, 108]. However, almost all of them are based on experiments, which always have threats to validity. There ideally should be some theoretical guidelines that can help researchers select the right predictive models according to the characteristics of the specific SE tasks.

The selection of classifiers is another possible source of bias. Given the variety of available learning algorithms, there are still others that could have been considered. An appropriate selection might be guided by the aim of finding a meaningful balance ,or trade-off, between established techniques and novel approaches.

8.2 Opportunities

8.2.1 Opportunity 1: Leveraging the power of big data. The scale of available SE data has rapidly become larger and larger. A growing number of studies tend to be large-scale. The biggest advantage of big data is that it generally leads to robust predictive models, improving the practicability and generality. If a study only uses several small datasets to build predictive models it may not be generalisable and its results may have serious threats to validity. On the contrary, results achieved using big data are usually more convincing and more likely to have generality. In addition, deep learning techniques usually require a large amount of data to learn their predictive models, since there are many more parameters to learn than when using traditional predictive models. With the trend to using big data, reducing the time and space cost for data computation can also become a challenging issue.

8.2.2 Opportunity 2: Neural network based predictive models. Studies that apply predictive models to software engineering tasks occurred over several decades. Although researchers have achieved significant improvement from simple linear regression to complex ensemble learning, they have been bottlenecks for addressing many software engineering tasks. Recently, deep learning, as an advanced predictive model, has become more and more popular. Some studies that tried to leverage deep learning for software engineering tasks have achieved even better performance than all other start-of-the-art techniques.

The biggest advantage of deep learning is that it can automatically generate more expressive features that are better for learning predictive models, which cant be done when using traditional predictive models [19]. As mentioned above, many software engineering tasks face the feature extraction challenge, i.e., either it is hard to manually infer proper features or there are too many features needing to be carefully selected for building their predictive models. Therefore, leveraging deep learning for various software engineering tasks to improve their performance is a promising direction.

A well-developed predictive model often requires training with a great number of instances. However, there is not enough data to build such models in certain research domains. This problem can be solved by leveraging a novel deep learning technique referred to as meta-learning [25] that simulates humans' behavior when learning new concepts and skills much faster and efficiently with many fewer data instances. So far, few studies in software engineering have adopted a good meta-learning model capable of well adapting or generalizing to new tasks and new environments that have never been encountered during training time. Therefore, in addition to constructing more large scale datasets, meta-learning could be adopted as another significant solution for training well-performed machine learning models with a few training examples. This technique has been verified in other fields such as image classification and robots.

8.2.3 Opportunity 3: Assessment and selection of predictive models. Clearly, computational efficiency and transparency are desirable features of candidate classifiers and appear to be a promising area for future research to formalize these concepts, e.g., by developing a multidimensional classifier assessment system [58].

Consequently, the assessment and selection of a predictive model should not be based on predictive accuracy alone but should be comprised of several additional criteria. These include computational efficiency, ease of use, and especially comprehensibility. Comprehensible models reveal the nature of the detected relationships and help improve our overall understanding of software failures and their sources, which, in turn, may enable the development of novel predictors in various research directions.

Efforts to design new software metrics and other explanatory variables appear to be a particularly promising area for future research. They have the potential to achieve general accuracy improvements across all types of classifiers.

8.2.4 Opportunity 4: predictive models in specific research domains. Predictive models in software requirements and testing: According to our analysis we found that predictive models were rarely applied in certain domains, such as software requirements, software design and implementation. There are many different practical problems that are able to be tackled by leveraging predictive models in these two topics. We believe that researchers can try to apply state-of-the-art models on classification, prioritization and prediction tasks in both software requirements and testing.

The usability analysis of prediction tasks: There is a significant difference between the accuracy and usefulness of predictions across projects. Projects with a very low percentage of positives and a very high number of classes are intuitively hard to predict for most approaches. However, this task is even harder for humans to perform manually. User studies are therefore needed to provide insights into when and where automated techniques are useful to humans, expecting to find a break even point where automation becomes more accurate than fully manual techniques.

Another potential direction is understanding the utility of change impact analysis. The ability to predict which classes are impacted by a new requirement can potentially support a range of tasks including refactoring decisions, defect prediction, and effort estimation [23]. An increased understanding of such tasks could enable developers to build more effective prediction algorithms.

Predictive models in identifying different defects: As mentioned above, a significant number of studies have applied predictive models to identify various types of defects, such as logical defects, syntax defects, interface defects, security defects, and performance defects. However, there are no studies that investigate on which defect types the existing defect predictive models work the best. In future work, researchers can conduct studies to investigate the performance of different categories of predictive models on predicting different types of defects. We may also add specific contexts when performing experiments.

9 CONCLUSION

We wanted to analyse the use of predictive models in software engineering. A comprehensive analysis was performed to answer our defined research questions using a systematic review covering 145 selected papers published between 2009 and 2019 in top SE venues. These studies were identified by following a systematic series of steps and assessing the quality of the studies. Our key findings from this study are summarized below:

- The cumulative number of predictive model related studies shows an increasing trend over the last decade, and most of the selected primary studies focus on proposing novel approaches.
- We found 52 different predictive models were employed in software engineering tasks. These models can be classified into three categories – base learners, ensemble learners and deep learners.
- Logistic Regression and Naive Bayes are the most widely used learning techniques to build predictive models for SE tasks to date. Several machine learning models are also popular models for addressing specific problems, including SVM and decision trees.
- Recall, precision, and F-measure are the most frequently used performance metrics for evaluating the effectiveness of predictive models.

- The application of deep learning techniques has seen an increasing trend in recent years. CNN and RNN are the two most popular deep learning models used in the studies for various SE tasks.
- Most predictive models have been applied to the defect prediction task and 7 of 11 research topics where predictive models are frequently used belong to software maintenance and quality assessment domains.
- We identified a number of challenges when using these predictive models for SE tasks, including issues with dataset quality, feature and evaluation metric selection, and lack of model selection guidelines.

The main objective of this survey was to analyze and classify the use of existing predictive models and their related studies in SE to date. A range of future studies should include more applications of predictive models to requirements engineering and testing, better assessment and selection guidelines for different models for different SE tasks and use of meta-learning with neural network-based models.

REFERENCES

- [1] Zahra Shakeri Hossein Abad, Vincenzo Gervasi, Didar Zowghi, and Behrouz H Far. 2019. Supporting analysts by dynamic extraction and classification of requirements-related knowledge. In *ICSE*. IEEE, 442–453.
- [2] Md Ahasanuzzaman, Muhammad Asaduzzaman, Chanchal K Roy, and Kevin A Schneider. 2020. CAPS: a supervised technique for classifying Stack Overflow posts concerning API issues. *EMSE* 25, 2 (2020), 1493–1532.
- [3] Preethu Rose Anish, Balaji Balasubramaniam, Abhishek Sainani, Jane Cleland-Huang, Maya Daneva, Roel J Wieringa, and Smita Ghaisas. 2016. Probing for requirements knowledge to stimulate architectural thinking. In *ICSE*. 843–854.
- [4] John Anvik and Gail C Murphy. 2011. Reducing the effort of bug report triage: Recommenders for development-oriented decisions. *TOSEM* 20, 3 (2011), 1–35.
- [5] Vitalii Avdiienko, Konstantin Kuznetsov, Alessandra Gorla, Andreas Zeller, Steven Arzt, Siegfried Rasthofer, and Eric Bodden. 2015. Mining apps for abnormal usage of sensitive data. In *FSE*, Vol. 1. IEEE, 426–436.
- [6] Alberto Bacchelli, Tommaso Dal Sasso, Marco D’Ambros, and Michele Lanza. 2012. Content classification of development emails. In *ICSE*. IEEE, 375–385.
- [7] Christian Bird, Adrian Bachmann, Eirik Aune, John Duffy, Abraham Bernstein, Vladimir Filkov, and Premkumar Devanbu. 2009. Fair and balanced? bias in bug-fix datasets. In *ESEC/FSE*. 121–130.
- [8] Kelly Blincoe, Giuseppe Valetto, and Daniela Damian. 2013. Do all task dependencies require coordination? the role of task properties in identifying critical coordination needs in software projects. In *FSE*. 213–223.
- [9] Pierre Bourque, Richard E Fairley, et al. 2014. *Guide to the software engineering body of knowledge (SWEBOK (R)): Version 3.0*. IEEE.
- [10] Jose Cambronero, Hongyu Li, Seohyun Kim, Koushik Sen, and Satish Chandra. 2019. When deep learning met code search. In *ESEC/FSE*. 964–974.
- [11] Gerardo Canfora, Massimiliano Di Penta, Rocco Oliveto, and Sebastiano Panichella. 2012. Who is going to mentor newcomers in open source projects?. In *FSE*. 1–11.
- [12] Mahinthan Chandramohan, Hee Beng Kuan Tan, Lionel C Briand, Lwin Khin Shar, and Bindu Madhavi Padmanabhuni. 2013. A scalable approach for malware detection through bounded feature space behavior modeling. In *ASE*. IEEE, 312–322.
- [13] Ning Chen, Steven CH Hoi, and Xiaokui Xiao. 2011. Software process evaluation: A machine learning approach. In *ASE*. IEEE, 333–342.
- [14] Xin Chen, He Jiang, Xiaochen Li, Liming Nie, Dongjin Yu, Tieke He, and Zhenyu Chen. 2020. A systemic framework for crowdsourced test report quality assessment. *EMSE* (2020), 1–37.
- [15] Zhenpeng Chen, Yanbin Cao, Xuan Lu, Qiaozhu Mei, and Xuanzhe Liu. 2019. SEntiMoji: an emoji-powered learning approach for sentiment analysis in software engineering. In *ESEC/FSE*. 841–852.
- [16] Morakot Choetkiertikul, Hoa Khanh Dam, Truyen Tran, and Aditya Ghose. 2015. Predicting delays in software projects using networked classification (t). In *ASE*. IEEE, 353–364.
- [17] Morakot Choetkiertikul, Hoa Khanh Dam, Truyen Tran, Trang Pham, Aditya Ghose, and Tim Menzies. 2018. A deep learning model for estimating story points. *TSE* 45, 7 (2018), 637–656.
- [18] Domenico Cotroneo, Roberto Pietrantuono, and Stefano Russo. 2013. A learning-based method for combining testing techniques. In *ICSE*. IEEE, 142–151.
- [19] Hoa Khanh Dam, Truyen Tran, Trang Thi Minh Pham, Shien Wee Ng, John Grundy, and Aditya Ghose. 2018. Automatic feature learning for predicting vulnerable software components. *TSE* (2018).
- [20] Karel Dejaeger, Wouter Verbeke, David Martens, and Bart Baesens. 2011. Data mining techniques for software effort estimation: a comparative study. *TSE* 38, 2 (2011), 375–397.

- [21] Andrea Di Sorbo, Sebastiano Panichella, Corrado A Visaggio, Massimiliano Di Penta, Gerardo Canfora, and Harald C Gall. 2015. Development emails content analyzer: Intention mining in developer discussions (T). In *ASE*. IEEE, 12–23.
- [22] Davide Falessi, Massimiliano Di Penta, Gerardo Canfora, and Giovanni Cantone. 2017. Estimating the number of remaining links in traceability recovery. *EMSE* 22, 3 (2017), 996–1027.
- [23] Davide Falessi, Justin Roll, Jin LC Guo, and Jane Cleland-Huang. 2018. Leveraging Historical Associations between Requirements and Source Code to Identify Impacted Classes. *TSE* (2018).
- [24] Yuanrui Fan, Xin Xia, Daniel Alencar da Costa, David Lo, Ahmed E Hassan, and Shanping Li. 2019. The Impact of Changes Misabeled by SZZ on Just-in-Time Defect Prediction. *TSE* (2019).
- [25] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400* (2017).
- [26] Francesca Arcelli Fontana, Mika V Mäntylä, Marco Zanoni, and Alessandro Marino. 2016. Comparing and experimenting machine learning techniques for code smell detection. *EMSE* 21, 3 (2016), 1143–1191.
- [27] Thomas Fritz, Andrew Begel, Sebastian C Müller, Serap Yigit-Elliott, and Manuela Züger. 2014. Using psycho-physiological measures to assess task difficulty in software development. In *ICSE*. 402–413.
- [28] Baljinder Ghotra, Shane McIntosh, and Ahmed E Hassan. 2015. Revisiting the impact of classification techniques on the performance of defect prediction models. In *ICSE*, Vol. 1. IEEE, 789–800.
- [29] Nils Göde and Rainer Koschke. 2011. Frequency and risks of changes to clones. In *ICSE*. 311–320.
- [30] Lina Gong, Shujuan Jiang, Rongcun Wang, and Li Jiang. 2019. Empirical Evaluation of the Impact of Class Overlap on Software Defect Prediction. In *ASE*. IEEE, 698–709.
- [31] Raghuram Gopalakrishnan, Palak Sharma, Mehdi Mirakhorli, and Matthias Galster. 2017. Can latent topics in source code predict missing architectural tactics?. In *ICSE*. IEEE, 15–26.
- [32] Philip J Guo, Thomas Zimmermann, Nachiappan Nagappan, and Brendan Murphy. 2010. Characterizing and predicting which bugs get fixed: an empirical study of Microsoft Windows. In *ICSE*. 495–504.
- [33] Andrew Habib and Michael Pradel. 2018. Is this class thread-safe? inferring documentation using graph-based learning. In *ASE*. 41–52.
- [34] Sonia Haiduc, Gabriele Bavota, Andrian Marcus, Rocco Oliveto, Andrea De Lucia, and Tim Menzies. 2013. Automatic query reformulations for text retrieval in software engineering. In *ICSE*. IEEE, 842–851.
- [35] Sonia Haiduc, Gabriele Bavota, Rocco Oliveto, Andrian Marcus, and Andrea De Lucia. 2012. Evaluating the specificity of text retrieval queries to support software engineering tasks. In *ICSE*. IEEE, 1273–1276.
- [36] Kim Herzig, Sascha Just, and Andreas Zeller. 2016. The impact of tangled code changes on defect prediction models. *EMSE* 21, 2 (2016), 303–336.
- [37] Eric Horton and Chris Parnin. 2019. V2: Fast Detection of Configuration Drift in Python. In *ASE*. IEEE, 477–488.
- [38] Seyedrebar Hosseini, Burak Turhan, and Dimuthu Gunarathna. 2017. A systematic literature review and meta-analysis on cross project defect prediction. *TSE* 45, 2 (2017), 111–147.
- [39] Seyedrebar Hosseini, Burak Turhan, and Mika Mäntylä. 2018. A benchmark study on the effectiveness of search-based data selection and feature selection for cross project defect prediction. *IST* 95 (2018), 296–312.
- [40] Qiao Huang, Emad Shihab, Xin Xia, David Lo, and Shanping Li. 2018. Identifying self-admitted technical debt in open source projects using text mining. *EMSE* 23, 1 (2018), 418–451.
- [41] Rubing Huang, Weifeng Sun, Yinyin Xu, Haibo Chen, Dave Towey, and Xin Xia. 2019. A Survey on Adaptive Random Testing. *TSE* (2019).
- [42] Gaehul Jeong, Sunghun Kim, and Thomas Zimmermann. 2009. Improving bug triage with bug tossing graphs. In *ESEC/FSE*. 111–120.
- [43] Nishant Jha and Anas Mahmoud. 2018. Using frame semantics for classifying and summarizing application store reviews. *EMSE* 23, 6 (2018), 3734–3767.
- [44] Xiaoyuan Jing, Fei Wu, Xiwei Dong, Fumin Qi, and Baowen Xu. 2015. Heterogeneous cross-company defect prediction by unified metric representation and CCA-based transfer learning. In *FSE*. 496–507.
- [45] Xiao-Yuan Jing, Fei Wu, Xiwei Dong, and Baowen Xu. 2016. An improved SDA based defect prediction framework for both within-project and cross-project class-imbalance problems. *TSE* 43, 4 (2016), 321–339.
- [46] Xiao-Yuan Jing, Shi Ying, Zhi-Wu Zhang, Shan-Shan Wu, and Jin Liu. 2014. Dictionary learning based software defect prediction. In *ICSE*. 414–423.
- [47] Leif Jonsson, Markus Borg, David Broman, Kristian Sandahl, Sigrid Eldh, and Per Runeson. 2016. Automated bug assignment: Ensemble-based machine learning in large scale industrial contexts. *EMSE* 21, 4 (2016), 1533–1578.
- [48] Magne Jørgensen. 2014. What we do and don't know about software development effort estimation. *IEEE software* 31, 2 (2014), 37–40.
- [49] Yasutaka Kamei, Takafumi Fukushima, Shane McIntosh, Kazuhiro Yamashita, Naoyasu Ubayashi, and Ahmed E Hassan. 2016. Studying just-in-time defect prediction using cross-project models. *EMSE* 21, 5 (2016), 2072–2106.
- [50] Yasutaka Kamei, Emad Shihab, Bram Adams, Ahmed E Hassan, Audris Mockus, Anand Sinha, and Naoyasu Ubayashi. 2012. A large-scale empirical study of just-in-time quality assurance. *TSE* 39, 6 (2012), 757–773.

- [51] Staffs Keele et al. 2007. *Guidelines for performing systematic literature reviews in software engineering*. Technical Report. Technical report, Ver. 2.3 EBSE Technical Report. EBSE.
- [52] Dongsun Kim, Yida Tao, Sunghun Kim, and Andreas Zeller. 2013. Where should we fix this bug? a two-phase recommendation model. *TSE* 39, 11 (2013), 1597–1610.
- [53] Sunghun Kim, E James Whitehead Jr, and Yi Zhang. 2008. Classifying software changes: Clean or buggy? *TSE* 34, 2 (2008), 181–196.
- [54] Sunghun Kim, Hongyu Zhang, Rongxin Wu, and Liang Gong. 2011. Dealing with noise in defect prediction. In *ICSE*. IEEE, 481–490.
- [55] Barbara A Kitchenham, Tore Dyba, and Magne Jorgensen. 2004. Evidence-based software engineering. In *ICSE*. IEEE, 273–281.
- [56] SB Kotsiantis, Dimitris Kanellopoulos, and PE Pintelas. 2006. Data preprocessing for supervised learning. *International Journal of Computer Science* 1, 2 (2006), 111–117.
- [57] Taek Lee, Jaechang Nam, DongGyun Han, Sunghun Kim, and Hoh Peter In. 2011. Micro interaction metrics for defect prediction. In *FSE*. 311–321.
- [58] Stefan Lessmann, Bart Baesens, Christophe Mues, and Swantje Pietsch. 2008. Benchmarking classification models for software defect prediction: A proposed framework and novel findings. *TSE* 34, 4 (2008), 485–496.
- [59] Heng Li, Weiyl Shang, Ying Zou, and Ahmed E Hassan. 2017. Towards just-in-time suggestions for log changes. *EMSE* 22, 4 (2017), 1831–1865.
- [60] Bin Lin, Fiorella Zampetti, Gabriele Bavota, Massimiliano Di Penta, and Michele Lanza. 2019. Pattern-based mining of opinions in Q&A websites. In *ICSE*. IEEE, 548–559.
- [61] Bin Lin, Fiorella Zampetti, Gabriele Bavota, Massimiliano Di Penta, Michele Lanza, and Rocco Oliveto. 2018. Sentiment Analysis for So ware Engineering: How Far Can We Go? (2018).
- [62] Dayi Lin, Cor-Paul Bezemer, and Ahmed E Hassan. 2019. Identifying gameplay videos that exhibit bugs in computer games. *EMSE* 24, 6 (2019), 4006–4033.
- [63] Hui Liu, Zhifeng Xu, and Yanzhen Zou. 2018. Deep learning based feature envy detection. In *ASE*. 385–396.
- [64] Pei Liu, Li Li, Yanjie Zhao, Xiaoyu Sun, and John Grundy. 2018. AndroZooOpen: Collecting Large-scale Open Source Android Apps for the Research Community. *Star* 1, 800 (2018), 1300.
- [65] Tianming Liu, Haoyu Wang, Li Li, Guangdong Bai, Yao Guo, and Guoai Xu. 2019. DaPanda: Detecting Aggressive Push Notifications in Android Apps. In *ASE*. IEEE, 66–78.
- [66] Xiaoyu Liu, LiGuo Huang, Jidong Ge, and Vincent Ng. 2019. Predicting Licenses for Changed Source Code. In *ASE*. IEEE, 686–697.
- [67] Lucia, David Lo, Lingxiao Jiang, Aditya Budi, et al. 2012. Active refinement of clone anomaly reports. In *ICSE*. IEEE, 397–407.
- [68] Mateusz Machalica, Alex Samykin, Meredith Porth, and Satish Chandra. 2019. Predictive test selection. In *ICSE SEIP*. IEEE / ACM, 91–100.
- [69] Chandra Maddila, Chetan Bansal, and Nachiappan Nagappan. 2019. Predicting pull request completion time: a case study on large scale cloud services. In *ESEC/FSE*. 874–882.
- [70] Onkar Malgonde and Kaushal Chari. 2019. An ensemble-based model for predicting agile software development effort. *EMSE* 24, 2 (2019), 1017–1055.
- [71] Ruchika Malhotra and Megha Khanna. 2017. An empirical study for software change prediction using imbalanced data. *EMSE* 22, 6 (2017), 2806–2851.
- [72] Haroon Malik, Hadi Hemmati, and Ahmed E Hassan. 2013. Automatic detection of performance deviations in the load testing of large scale systems. In *ICSE*. IEEE, 1012–1021.
- [73] Daniel Martens and Walid Maalej. 2019. Towards understanding and detecting fake reviews in app stores. *EMSE* 24, 6 (2019), 3316–3355.
- [74] Andrew Meneely, Pete Rotella, and Laurie Williams. 2011. Does adding manpower also affect quality? an empirical, longitudinal analysis. In *ESEC/FSE*. 81–90.
- [75] Tom Mens and Tom Tourwé. 2004. A survey of software refactoring. *TSE* 30, 2 (2004), 126–139.
- [76] Tim Menzies, Ye Yang, George Mathew, Barry Boehm, and Jaius Hihn. 2017. Negative results for software effort estimation. *EMSE* 22, 5 (2017), 2658–2683.
- [77] Amin Milani Fard, Mehdi Mirzaaghaei, and Ali Mesbah. 2014. Leveraging existing tests in automated test generation for web applications. In *ASE*. 67–78.
- [78] Chris Mills, Gabriele Bavota, Sonia Haiduc, Rocco Oliveto, Andrian Marcus, and Andrea De Lucia. 2017. Predicting query quality for applications of text retrieval to software engineering tasks. *TOSEM* 26, 1 (2017), 1–45.
- [79] Mehdi Mirakhorli and Jane Cleland-Huang. 2015. Detecting, tracing, and monitoring architectural tactics in code. *TSE* 42, 3 (2015), 205–220.
- [80] Mehdi Mirakhorli, Yonghee Shin, Jane Cleland-Huang, and Murat Cinar. 2012. A tactic-centric approach for automating traceability of quality concerns. In *ICSE*. IEEE, 639–649.
- [81] Ayse Tosun Misirli, Emad Shihab, and Yasukata Kamei. 2016. Studying high impact fix-inducing changes. *EMSE* 21, 2 (2016), 605–641.

- [82] Nikolaos Mittas and Lefteris Angelis. 2012. Ranking and clustering software cost estimation models through a multiple comparisons algorithm. *TSE* 39, 4 (2012), 537–551.
- [83] Audris Mockus. 2009. Amassing and indexing a large sample of version control systems: Towards the census of public source code history. In *MSR*. IEEE, 11–20.
- [84] Toshiki Mori and Naoshi Uchihira. 2019. Balancing the trade-off between accuracy and interpretability in software defect prediction. *EMSE* 24, 2 (2019), 779–825.
- [85] Golam Mostaeen, Jeffrey Svajlenko, Banani Roy, Chanchal K Roy, and Kevin A Schneider. 2019. CloneCognition: machine learning based code clone validation tool. In *ESEC/FSE*. 1105–1109.
- [86] Sebastian C Müller and Thomas Fritz. 2015. Stuck and frustrated or in flow and happy: sensing developers’ emotions and progress. In *ICSE*, Vol. 1. IEEE, 688–699.
- [87] Sebastian C Müller and Thomas Fritz. 2016. Using (bio) metrics to predict code quality online. In *ICSE*. IEEE, 452–463.
- [88] Nuthan Munaiah, Steven Kroh, Craig Cabrey, and Meiyappan Nagappan. 2017. Curating GitHub for engineered software projects. *EMSE* 22, 6 (2017), 3219–3253.
- [89] Pradeep K Murukannaiah and Munindar P Singh. 2015. Platys: An active learning framework for place-aware application development and its evaluation. *TOSEM* 24, 3 (2015), 1–32.
- [90] Kawser Wazed Nafi, Tonny Shekha Kar, Banani Roy, Chanchal K Roy, and Kevin A Schneider. 2019. CLCDSA: cross language code clone detection using syntactical features and API documentation. In *ASE*. IEEE, 1026–1037.
- [91] Nachiappan Nagappan and Thomas Ball. 2005. Use of relative code churn measures to predict system defect density. In *ICSE*. 284–292.
- [92] Vivek Nair, Tim Menzies, Norbert Siegmund, and Sven Apel. 2017. Using bad learners to find good configurations. In *FSE*. 257–267.
- [93] Daye Nam, Amber Horvath, Andrew Macvean, Brad Myers, and Bogdan Vasilescu. 2019. Marble: Mining for boilerplate code to identify API usability problems. In *ASE*. IEEE, 615–627.
- [94] Jaechang Nam, Wei Fu, Sunghun Kim, Tim Menzies, and Lin Tan. 2017. Heterogeneous defect prediction. *TSE* 44, 9 (2017), 874–896.
- [95] Jaechang Nam and Sunghun Kim. 2015. Clami: Defect prediction on unlabeled datasets (t). In *ASE*. IEEE, 452–463.
- [96] Jaechang Nam, Sinno Jialin Pan, and Sunghun Kim. 2013. Transfer defect learning. In *ICSE*. IEEE, 382–391.
- [97] Ning Nan and Donald E Harter. 2009. Impact of budget and schedule pressure on software development cycle time and effort. *TSE* 35, 5 (2009), 624–637.
- [98] Roberto Natella, Domenico Cotroneo, Joao A Duraes, and Henrique S Madeira. 2012. On fault representativeness of software fault injection. *TSE* 39, 1 (2012), 80–96.
- [99] Haoran Niu, Iman Keivanloo, and Ying Zou. 2017. Learning to rank code examples for code search engines. *EMSE* 22, 1 (2017), 259–291.
- [100] Burcu Ozelik and Cemal Yilmaz. 2015. Seer: a lightweight online failure prediction approach. *TSE* 42, 1 (2015), 26–46.
- [101] Rohan Padhye, Senthil Mani, and Vibha Singhal Sinha. 2014. NeedFeed: taming change notifications by modeling code relevance. In *ASE*. 665–676.
- [102] Fabio Palomba, Damian Andrew Andrew Tamburri, Francesca Arcelli Fontana, Rocco Oliveto, Andy Zaidman, and Alexander Serebrenik. 2018. Beyond technical aspects: How do community smells influence the intensity of code smells? *TSE* (2018).
- [103] Luca Pascarella, Magiel Bruntink, and Alberto Bacchelli. 2019. Classifying code comments in Java software systems. *EMSE* 24, 3 (2019), 1499–1537.
- [104] Anna Perini, Angelo Susi, and Paolo Avesani. 2012. A machine learning approach to software requirements prioritization. *TSE* 39, 4 (2012), 445–461.
- [105] Fayola Peters and Tim Menzies. 2012. Privacy and utility for defect prediction: Experiments with morph. In *ICSE*. IEEE, 189–199.
- [106] Fayola Peters, Tim Menzies, Liang Gong, and Hongyu Zhang. 2013. Balancing privacy and utility in cross-company defect prediction. *TSE* 39, 8 (2013), 1054–1068.
- [107] Fayola Peters, Tim Menzies, and Lucas Layman. 2015. LACE2: Better privacy-preserving data sharing for cross project defect prediction. In *ICSE*, Vol. 1. IEEE, 801–811.
- [108] Fayola Peters, Thein Tun, Yijun Yu, and Bashar Nuseibeh. 2017. Text filtering and ranking for security bug report prediction. *TSE* (2017).
- [109] Kai Petersen, Sairam Vakkalanka, and Ludwik Kuzniarz. 2015. Guidelines for conducting systematic mapping studies in software engineering: An update. *Information and Software Technology* 64 (2015), 1–18.
- [110] Gayane Petrosyan, Martin P Robillard, and Renato De Mori. 2015. Discovering information explaining API types using text classification. In *ICSE*, Vol. 1. IEEE, 869–879.
- [111] Andy Podgurski, David Leon, Patrick Francis, Wes Masri, Melinda Minch, Jiayang Sun, and Bin Wang. 2003. Automated support for classifying software failure reports. In *ICSE*. IEEE, 465–475.
- [112] Gede Artha Azriadi Prana, Christoph Treude, Ferdian Thung, Thushari Atapattu, and David Lo. 2019. Categorizing the content of GitHub README files. *EMSE* 24, 3 (2019), 1296–1327.
- [113] Foyzur Rahman and Premkumar Devanbu. 2013. How, and why, process metrics are better. In *ICSE*. IEEE, 432–441.

- [114] Foyzur Rahman, Sameer Khatri, Earl T Barr, and Premkumar Devanbu. 2014. Comparing static bug finders and statistical prediction. In *ICSE*. 424–434.
- [115] Foyzur Rahman, Daryl Posnett, Israel Herraiz, and Premkumar Devanbu. 2013. Sample size vs. bias in defect prediction. In *FSE*. 147–157.
- [116] Inayat Rehman, Mehdi Mirakhorli, Meiyappan Nagappan, Azat Aralbay Uulu, and Matthew Thornton. 2018. Roles and impacts of hands-on software architects in five industrial case studies. In *ICSE*. 117–127.
- [117] Barbara Russo, Giancarlo Succi, and Witold Pedrycz. 2015. Mining system logs to learn error predictors: a case study of a telemetry system. *EMSE* 20, 4 (2015), 879–927.
- [118] Vaibhav Saini, Farima Farmahinifarahani, Yadong Lu, Pierre Baldi, and Cristina V Lopes. 2018. Oreo: Detection of clones in the twilight zone. In *ESEC/FSE*. 354–365.
- [119] Felix Salfner, Maren Lenk, and Mirosław Malek. 2010. A survey of online failure prediction methods. *ACM Computing Surveys (CSUR)* 42, 3 (2010), 1–42.
- [120] Riccardo Scandariato, James Walden, Aram Hovsepyan, and Wouter Joosen. 2014. Predicting vulnerable software components via text mining. *TSE* 40, 10 (2014), 993–1006.
- [121] Jürgen Schmidhuber. 2015. Deep learning in neural networks: An overview. *Neural networks* 61 (2015), 85–117.
- [122] Linda Shafer. 2014. (CSDP) Software Engineering Professional Practice. (2014).
- [123] Lwin Khin Shar, Hee Beng Kuan Tan, and Lionel C Briand. 2013. Mining SQL injection and cross site scripting vulnerabilities using hybrid program analysis. In *ICSE*. IEEE, 642–651.
- [124] Abhishek Sharma, Yuan Tian, Agus Sulistya, Dinusha Wijedasa, and David Lo. 2018. Recommending who to follow in the software engineering twitter space. *TOSEM* 27, 4 (2018), 1–33.
- [125] Yonghee Shin, Andrew Meneely, Laurie Williams, and Jason A Osborne. 2010. Evaluating complexity, code churn, and developer activity metrics as indicators of software vulnerabilities. *TSE* 37, 6 (2010), 772–787.
- [126] Shivkumar Shivaji, E James Whitehead, Ram Akella, and Sunghun Kim. 2012. Reducing features to improve code change-based bug prediction. *TSE* 39, 4 (2012), 552–569.
- [127] Charles Song, Adam Porter, and Jeffrey S Foster. 2013. iTree: efficiently discovering high-coverage configurations using interaction trees. *TSE* 40, 3 (2013), 251–265.
- [128] Charles Song, Adam A. Porter, and Jeffrey S. Foster. 2012. iTree: Efficiently discovering high-coverage configurations using interaction trees. In *ICSE*. IEEE, 903–913.
- [129] Liyan Song, Leandro L Minku, and Xin Yao. 2018. A novel automated approach for software effort estimation based on data augmentation. In *FSE*. 468–479.
- [130] Qinqin Song, Zihan Jia, Martin Shepperd, Shi Ying, and Jin Liu. 2010. A general software defect-proneness prediction framework. *TSE* 37, 3 (2010), 356–370.
- [131] Chengnian Sun, David Lo, Siau-Cheng Khoo, and Jing Jiang. 2011. Towards more accurate retrieval of duplicate bug reports. In *ASE*. IEEE, 253–262.
- [132] Chengnian Sun, David Lo, Xiaoyin Wang, Jing Jiang, and Siau-Cheng Khoo. 2010. A discriminative model approach for accurate duplicate bug report retrieval. In *ICSE*. 45–54.
- [133] Hao Sun, Xiangyu Zhang, Yunhui Zheng, and Qingkai Zeng. 2016. Inteq: recognizing benign integer overflows via equivalence checking across multiple precisions. In *ICSE*. 1051–1062.
- [134] Chakkrit Tantithamthavorn, Ahmed E Hassan, and Kenichi Matsumoto. 2018. The impact of class rebalancing techniques on the performance and interpretation of defect prediction models. *TSE* (2018).
- [135] Chakkrit Tantithamthavorn, Shane McIntosh, Ahmed E Hassan, Akinori Ihara, and Kenichi Matsumoto. 2015. The impact of mislabelling on the performance and interpretation of defect prediction models. In *ICSE*, Vol. 1. IEEE, 812–823.
- [136] Chakkrit Tantithamthavorn, Shane McIntosh, Ahmed E Hassan, and Kenichi Matsumoto. 2016. Automated parameter optimization of classification techniques for defect prediction models. In *ICSE*. 321–332.
- [137] Christopher Theisen, Kim Herzig, Patrick Morrison, Brendan Murphy, and Laurie Williams. 2015. Approximating attack surfaces with stack traces. In *ICSE*, Vol. 2. IEEE, 199–208.
- [138] Stephen W Thomas, Meiyappan Nagappan, Dorothea Blostein, and Ahmed E Hassan. 2013. The impact of classifier configuration and classifier combination on bug localization. *TSE* 39, 10 (2013), 1427–1443.
- [139] Patanamon Thongtanunam, Weiyi Shang, and Ahmed E Hassan. 2019. Will this clone be short-lived? Towards a better understanding of the characteristics of short-lived clones. *EMSE* 24, 2 (2019), 937–972.
- [140] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. 2018. Deeptest: Automated testing of deep-neural-network-driven autonomous cars. In *ICSE*. 303–314.
- [141] Burak Turhan, Tim Menzies, Ayşe B Bener, and Justin Di Stefano. 2009. On the relative value of cross-company and within-company data for defect prediction. *EMSE* 14, 5 (2009), 540–578.

- [142] Christopher Vendome, Mario Linares-Vásquez, Gabriele Bavota, Massimiliano Di Penta, Daniel German, and Denys Poshyvanyk. 2017. Machine learning-based detection of open source license exceptions. In *ICSE*. IEEE, 118–129.
- [143] Junjie Wang, Song Wang, Qiang Cui, and Qing Wang. 2016. Local-based active classification of test report to assist crowdsourced testing. In *ASE*. 190–201.
- [144] Song Wang, Taiyue Liu, and Lin Tan. 2016. Automatically learning semantic features for defect prediction. In *ICSE*. IEEE, 297–308.
- [145] Xiaoyin Wang, Yingnong Dang, Lu Zhang, Dongmei Zhang, Erica Lan, and Hong Mei. 2012. Can I clone this piece of code here?. In *ICSE*. 170–179.
- [146] Xiaoyin Wang, Yingnong Dang, Lu Zhang, Dongmei Zhang, Erica Lan, and Hong Mei. 2014. Predicting consistency-maintenance requirement of code clones at copy-and-paste time. *TSE* 40, 8 (2014), 773–794.
- [147] Peter A Whigham, Caitlin A Owen, and Stephen G Macdonell. 2015. A baseline model for software effort estimation. *TOSEM* 24, 3 (2015), 1–11.
- [148] Andrew Wood, Paige Rodeghero, Ameer Armaly, and Collin McMillan. 2018. Detecting speech act types in developer question/answer conversations during bug repair. In *ESEC/FSE*. 491–502.
- [149] Bowen Xu, Deheng Ye, Zhenchang Xing, Xin Xia, Guibin Chen, and Shanping Li. 2016. Predicting semantically linkable knowledge in developer online forums via convolutional neural network. In *ASE*. IEEE, 51–62.
- [150] Hui Yang, Alistair Willis, Anne De Roeck, and Bashar Nuseibeh. 2010. Automatic detection of nocuous coordination ambiguities in natural language requirements. In *ASE*. 53–62.
- [151] Yibiao Yang, Yuming Zhou, Jinping Liu, Yangyang Zhao, Hongmin Lu, Lei Xu, Baowen Xu, and Hareton Leung. 2016. Effort-aware just-in-time defect prediction: simple unsupervised models could be better than supervised models. In *FSE*. 157–168.
- [152] Cemal Yilmaz and Adam Porter. 2010. Combining hardware and software instrumentation to classify program executions. In *FSE*. 67–76.
- [153] Tingting Yu, Wei Wen, Xue Han, and Jane Huffman Hayes. 2018. Conpredictor: concurrency defect prediction in real-world applications. *TSE* 45, 6 (2018), 558–575.
- [154] Zhongxing Yu, Chenggang Bai, and Kai-Yuan Cai. 2015. Does the failing test execute a single or multiple faults? An approach to classifying failing tests. In *ICSE*, Vol. 1. IEEE, 924–935.
- [155] Marcelo Serrano Zanetti, Ingo Scholtes, Claudio Juan Tessone, and Frank Schweitzer. 2013. Categorizing bugs with social networks: a case study on four open source software communities. In *ICSE*. IEEE, 1032–1041.
- [156] Feng Zhang, Quan Zheng, Ying Zou, and Ahmed E Hassan. 2016. Cross-project defect prediction using a connectivity-based unsupervised classifier. In *ICSE*. IEEE, 309–320.
- [157] Hongyu Zhang, Liang Gong, and Steve Versteege. 2013. Predicting bug-fixing time: an empirical study of commercial software projects. In *ICSE*. IEEE, 1042–1051.
- [158] Jie Zhang, Lingming Zhang, Mark Harman, Dan Hao, Yue Jia, and Lu Zhang. 2018. Predictive mutation testing. *TSE* 45, 9 (2018), 898–918.
- [159] Kevin Zhang. 2019. A Machine Learning Based Approach to Identify SQL Injection Vulnerabilities. In *ASE*. IEEE, 1286–1288.
- [160] Xu Zhang, Yong Xu, Qingwei Lin, Bo Qiao, Hongyu Zhang, Yingnong Dang, Chunyu Xie, Xinsheng Yang, Qian Cheng, Ze Li, et al. 2019. Robust log-based anomaly detection on unstable log data. In *ESEC/FSE*. 807–817.
- [161] Wujie Zheng, Haochuan Lu, Yangfan Zhou, Jianming Liang, Haibing Zheng, and Yuetang Deng. 2019. iFeedback: Exploiting User Feedback for Real-Time Issue Detection in Large-Scale Online Service Systems. In *ASE*. IEEE, 352–363.
- [162] Minghui Zhou and Audris Mockus. 2014. Who will stay in the floss community? modeling participants initial behavior. *TSE* 41, 1 (2014), 82–99.
- [163] Xiang Zhou, Xin Peng, Tao Xie, Jun Sun, Chao Ji, Dewei Liu, Qilin Xiang, and Chuan He. 2019. Latent error prediction and fault localization for microservice applications by learning from system trace logs. In *ESEC/FSE*. 683–694.
- [164] Yuming Zhou, Baowen Xu, Hareton Leung, and Lin Chen. 2014. An in-depth study of the potentially confounding effect of class size in fault prediction. *TOSEM* 23, 1 (2014), 1–51.
- [165] Yuming Zhou, Yibiao Yang, Hongmin Lu, Lin Chen, Yanhui Li, Yangyang Zhao, Junyan Qian, and Baowen Xu. 2018. How far we have progressed in the journey? an examination of cross-project defect prediction. *TOSEM* 27, 1 (2018), 1–51.
- [166] Zhi-Hua Zhou. 2009. Ensemble Learning. *Encyclopedia of biometrics* 1 (2009), 270–273.
- [167] Thomas Zimmermann, Nachiappan Nagappan, Harald Gall, Emanuel Giger, and Brendan Murphy. 2009. Cross-project defect prediction: a large scale experiment on data vs. domain vs. process. In *ESEC/FSE*. 91–100.
- [168] Thomas Zimmermann, Nachiappan Nagappan, Philip J Guo, and Brendan Murphy. 2012. Characterizing and predicting which bugs get reopened. In *ICSE*. IEEE, 1074–1083.