Data-Driven Decisions and Actions in Today's Software Development

Harald Gall, Carol Alexandru, Adelina Ciurumelea, Giovanni Grano,
Christoph Laaber, Sebastiano Panichella, Sebastian Proksch,
Gerald Schermann, Carmine Vassallo, and Jitong Zhao

1 Introduction

When software development is portrayed, it is often shown as a rather uncoordinated activity in which some genius programmers get together to hack a program until it magically starts working. While this picture may have been arguable decades ago, professional software development is a very structured activity. It has evolved into a complex process that is heavily relying on various kinds of data about the software itself, about its execution environment, but also about feedback from its users and the market. All these data are used to fuel the development for optimizing technical aspects and user experience. Today's processes fully integrate all phases of software engineering, from requirements to coding and testing, to integration and deployment, and to operations of the software system. Boundaries of development and operations have been reduced to the extent that both activities have been integrated and optimized leading to better performing software in the field.

Today's software development is very well tracked and recorded in various forms of data: requirements data, code and bug repositories, testing data and code, configurations and continuous integration data, deployment scripts, or data in social coding platforms such as StackOverflow. Additionally, runtime data (e.g., from executions in the cloud) are collected and fed back into the development process.

H. Gall · C. Alexandru · A. Ciurumelea · G. Grano · C. Laaber · S. Panichella · S. Proksch (✉) · G. Schermann · C. Vassallo

Department of Informatics, University of Zurich, Zurich, Switzerland

e-mail: gall@ifi.uzh.ch; alexandru@ifi.uzh.ch; ciurumelea@ifi.uzh.ch; grano@ifi.uzh.ch; laaber@ifi.uzh.ch; panichella@ifi.uzh.ch; proksch@ifi.uzh.ch; schermann@ifi.uzh.ch; vassallo@ifi.uzh.ch

J. Zhao

Tongji University, Shanghai, China

e-mail: 1410787@tongji.edu.cn

Software developers have to deal with a lot of infrastructure and data, be it static or dynamic data in various forms. Software experimentation provides further means to optimize feature roll-out for different user bases on a global scale. Consequently, process support has become even more essential given this mass of data and the interleaving processes.

From a business perspective, development workflows have been designed to allow the software creation to be planable by estimating efforts and monitoring progress. From a technical perspective, these workflows represent a safety net for the developers, because they both facilitate quality control and help to structure the work.

Inspired by the typical workflows of agile software development, Fig. 1 illustrates one typical release cycle in a simplified workflow. Each cycle starts from a list of requirements and leads to a shipped product, with several phases in between. We have analyzed these phases and have identified several challenges that developers are facing in these individual phases. In the following, we describe these challenges, the data to be used for decision making, and the actions that can be supported. For that we present opportunities to improve over the current state of the art by adapting results and tools that are proposed by recent research.

In Fig. 1, the first step in a release cycle is the planning phase, in which the scope of the next release is planned (1). This step is typically not a technical problem, but driven by business decisions. It is, therefore, considered separate and out of scope for our discussion.

Fig. 1 Release cycle

The core part of the release cycle is the implementation of the product itself (2). Modern software engineering has become less centered around the efficient implementation of algorithms, but more about a smart composition of existing components. Large numbers of frameworks and libraries exist that can be reused to solve a problem at hand. Reuse has several benefits like maintainability and maturity of the code base, but it requires the developer to constantly learn new application programming interfaces (APIs) to use these existing components effectively. Especially inexperienced developers or developers that switch projects are confronted with a huge technology stack that needs to be mastered, before they can contribute to the project. This overwhelming amount of information is not only very frustrating but also hinders productivity.

Heavily interwoven with the actual creation of source code is the creation of test code that double-checks that the implementation of the developer follows the expected behavior (3). Extensive testing is crucial to achieve a high software quality. Traditionally, developer either write small unit tests that test individual classes or resort to manual debugging to assert the expected behavior. Both approaches are suboptimal though, as important test scenarios might be missed. More recently, approaches appeared that can automatically generate test suits. Unfortunately, the quality of the generated tests is low and these suits are hard to maintain. In addition, some components are not easy to test automatically (e.g., user interfaces) and new ways of testing are required.

The most natural artifact that is created in the release cycle is source code. Both the implementation phase (2) and the testing phase (3) create source code as the primary artifact. One of the most important means for quality control in a project is a review of this source code before it can be considered finished. On top of this, source code is a way to communicate with other developers, and developers are required to understand source code on a daily basis. The high amount of source code is a big challenge for developers as it is often hard to find the important bits and pieces.

An important part of a modern release cycle is the existence of a build server that performs all integration tasks on every commit (4). Such a continuous integration/continuous deployment system asserts that the current code base can be used on different systems, for example, it would detect, if the developers forgot to add files to a commit or broke a configuration file. In addition to asserting the deployability of the product, the build server often performs several actions to enforce several quality attributes (e.g., test coverage, code style, code complexity). When build servers fail, they typically maintain a build log that reports about the individual steps of the build. Unfortunately, the cause of the build error is often not obvious. Developers need to understand the errors reported in a build to fix them in source code or through adapting the build configuration. The length of a build log and cryptic error messages take time to process and present a big challenge for developers.

After completing several iterations in this cycle, the updated product will eventually be deployed (5). This can happen in multiple ways; most commonly, a release of the product is created and published as an update to its users. For

long-running projects, the number of existing releases is large and projects that exercise continuous delivery principles excessively might create multiple releases per day. Analyzing the historic development of these releases is not only interesting for research but also for project managers that want to study telemetry data of their product. By monitoring trends about, for example, performance metrics, code complexity, or the number of dependencies, it is possible to identify degradations early and to make educated and data-driven decisions about technical details of the product.

A crosscutting concern that is valid across all phases is the review of changes (6). We already motivated the need for summarization techniques to support efficient reviews of source code reviews, but other use cases exist. In case of a deployment to app stores, any release can also be rated by users, which creates a valuable source of feedback for the developers. The vast numbers of ratings and reviews can be leveraged to better understand the requirements and sentiments of the target users. Rating and reviews can be used to decide about the direction of future development. However, to make the amount of data digestible by a human, it is required to summarize it to the most valuable information first.

In the following, we will dedicate one section to each phase of the release cycle. We describe current tools and approaches that add value to the process, describe the data-driven aspects and the current boundaries of feasibility, and sketch the next big idea that is about to impact the corresponding phase. While this chapter has a survey character, we will restrict the referenced works to existing approaches and to promising new results that could be put to action soon. We do, however, omit some visionary solutions for which the adoption cannot be expected in the near future.

2 Recommendation

Programmers extensively leverage application programming interfaces (APIs) to reuse code and unify programming experience when writing code. However, it is still a tough task for programmers to choose and utilize APIs from numerous libraries and frameworks. Even the most experienced programmers may encounter unfamiliar APIs and spend lots of time to learn how to use it. Furthermore, various barriers [68, 106] cause APIs to be hard to learn, such as insufficient examples and ambiguous documentation. As a result, it is a critical job for assisting programmers to learn APIs effectively and efficiently with less effort.

2.1 Code Example Recommendation Systems

To help programmers alleviate burden and better facilitate the usage of APIs, developers can be supported through code example recommendation systems [97]. Over the last decade, several of such systems have been proposed [95, 102, 108].

These systems propose programming information, such as which methods are most likely to call and how to invoke these methods. Traditional code example recommendation systems generally provide suggestions based on API usage data. Existing contributions can be organized in the categories according to the purpose of the detecting techniques.

Code search engines (CSEs) such as Krugle¹ and NerdyData² usually leverage text-oriented information-retrieval (IR) techniques to search in a large number of open-source repositories. But they don't provide mechanisms to rank the quality of the found code snippets and usually return too massive results for programmers to choose.

API usage example recommendation systems utilize API examples or API calls to recommend example code. These systems adopt various categories of API usage patterns and various techniques for inferring and clustering API usage algorithms, such as DBSCAN [109], k-nearest neighbor [16], BIDE [128], and RTM [83], or clustering algorithms like Canopy clustering [98]. The quality of the API usage examples found by these tools is derived from the overall quality of the code repositories they utilize and the selected mining algorithms.

Other tools use a semantic analysis approach to explore API usage obstacles through analyzing programmers' questions in Q&A website. Example Overflow [129] uses keyword search based on the Apache Lucene [76] library, which internally uses the term frequency-inverse document frequency (tf-idf) weight [127]. Using Q&A website as code repositories, systems would not be able to critically evaluate various snippets, and bugs may crop up for the examples that are not properly tested.

2.2 *Naturalness of Software*

The implementation of recommender systems, which aid developers in writing and maintaining code, has often employed machine learning and data mining approaches. The availability of a large and growing body of successful open-source projects and a recent hypothesis, "the naturalness of software", has opened the possibility of applying natural language processing techniques to source code. The hypothesis states that software, as a form of human communication, has statistical properties similar to the ones specific to natural language and that these can be exploited to build and improve software engineering tools [7].

Code suggestion is one of the most popular recommender systems and most used features of any modern IDE; it is typically implemented using manually defined syntactic and semantic completion rules derived from the programming language specification. Hindle et al. [60] observed that code corpora present a high degree

¹<http://www.krugle.com/>.

²<https://nerdydata.com/>.

of repetitiveness and they were able to exploit this property using a simple n-gram language model to enhance the code suggestion capabilities of Eclipse. Allamanis et al. [4] take advantage of the available open-source code online and learn a language model using a corpus 100 times larger than the previous work and improve their results showing that language models learned over source code, just like natural language, benefit significantly from more data. Tu et al. [120] analyze the limitation of previous models in capturing local regularities that are highly typical for human written programs and build a cache language model that further improves the code suggestion accuracy of previous work.

In [5] the authors tackle an interesting problem using statistical language models of source code: that of learning coding conventions from a code base. Adhering to coding conventions is an important practice of any successful and high-quality software project, as it strongly impacts readability and maintainability and it is often enforced by developers, but it currently lacks support in modern IDEs. Allamanis et al. [5] learn the coding style conventions specific to a software project through simple n-gram-based language models that are subsequently used to detect style violations of identifier naming and formatting and suggest improvements. One of the limits of this approach is that it can only suggest names that appear in the training set of the language models. While this is an adequate solution for local variable names, suggesting method and class names requires a more sophisticated approach. In [6], the authors experiment with a neural log-bilinear language model that is able to recommend neologism, names that do not appear in the training corpus, with promising results.

Natural language processing techniques applied to source code are extremely versatile: researchers have exploited them to evaluate code contributions to open-source projects and analyze whether they are likely to be accepted or not [59], improve reporting of compiler error messages [19], help developers find buggy code that is flagged as unnatural by language models [105], etc. Nevertheless, these techniques come with their own set of challenges. Natural language and source code have different characteristics which have to be taken into account when reusing approaches built and evaluated primarily for spoken and written language. A second problem is that a basic principle of software engineering, reusability, creates a data sparsity problem: it is rare to find multiple implementations of the same task in code, while it is quite common to find many news articles written about the same topic. While programming, developers often define new terms and compose them in novel ways; current NLP methods for natural language texts have been developed expecting that this is unlikely to occur. Another important issue is the evaluation of models trained on source code; there is a need of metrics adapted to source code and of existing benchmarks for researchers to compare their results. In spite of the existing limitations, there is a wide potential to apply natural language processing methods in a wide range of areas in software engineering and support developers in writing and maintaining code.

2.3 Evaluation

Independently of the technology that is used to build a recommendation system, be it a traditional recommender learned from examples or a recommender that is built on top of a language model, it will always remain a great challenge to evaluate the value of such a system. Traditionally, researchers have built benchmarks from the vast amounts of source code found in public repositories like GitHub. They would use the source code to learn models and validate the models on other examples. This approach has one significant drawback though: these benchmarks need to be considered as artificial. Previous work has shown that the history recorded in a repository is not representative for actual development [87], because it is incomplete. They find that a representative picture requires more fine-grained evolution data.

To close this mismatch, researchers have started to collect interaction data of developers directly in the IDE. The tools DFlow [79], FeedBaG [8, 101], or Epicea [39] are three examples of such systems that track developers during their day-to-day activities. The resulting datasets [100] present a unique opportunity to learn about patterns in developer behavior and to identify chances to improve their productivity in reoccurring tasks.

Examples of how to use such information are presented in several studies on the typical time budget of developers [9], frequently used commands in the IDE [85], or to find smells in interaction sequences [32]. As some of these trackers also capture source code changes, it is possible to use the interaction data as a ground truth for the evaluation of recommendation systems in software engineering. Prior work has shown that these realistic evaluations report different quality metrics for recommenders, when the evolving nature of source code is not reflected in the evaluation setup [99].

3 Testing

It is widely recognized that software testing is a crucial task in any successful software development process. Indeed, the overall testing cost has been estimated at being about half of the total development cost [12]. The definition of software testing involves several different kinds of activities and processes. In fact, various types of testing need to be performed in order to achieve different objectives and assess the qualities and the reliability of a software system. There are two main categories in software testing. On one side, *functional testing* assesses whether software behaves as intended. This category contains unit, integration, and user acceptance testing. On the other side, *nonfunctional testing* is concerned with program attributes like performance, security, or reliability.

Software testing is extensively handled in research; hence, we focus on two specific topics from both abovementioned categories. We start off by introducing concepts from automated unit test generation. Afterwards, performance testing in the form of software microbenchmarking is discussed.

3.1 Automated Unit Test Case Generation

Unit test is intended as a piece of code that automatically invokes a *unit* of work in a given system, checking assumptions about its behavior. To do that, inputs that exercise such units need to be defined. However, finding those inputs and writing test cases for a large system are extremely costly, difficult, and laborious tasks. An obvious response to this problem is to automate such a process as much as possible. Since the test case generation problem can be easily expressed as an optimization one [3], a tremendous amount of research has been conducted in applying metaheuristic algorithms (especially Genetics Algorithms (GA) [51]) to solve such a problem. The Search-Based Software Testing (SBST), an entirely new line of research, is the result of such a growing interest in the area [77].

The design of a search algorithm tailored to solve any optimization problem usually starts from the definition of the solution representation and the fitness function. In this context, a solution is represented by a set of test inputs. The fitness function is used to represent how *good* is a given solution for a coverage criterion. The most common one is the branch coverage criterion. A fitness function is composed mainly by two metrics: the *approach level* and the *branch distance*. The former expresses how far is the actual execution path from covering the target; the latter represents how far is the input data to change the Boolean value in the closer condition node to the target. Depending on how the targets are handled by the evolutionary algorithm, we can distinguish between single-target and multi-target approaches.

3.1.1 Single-Target Approaches

This class of algorithms has been the first one proposed in the literature as a search-based approach for test case generation [21, 42, 49]. A single-target strategy works as follows: (1) all the targets to hit are listed, (2) a single-objective search algorithm is used to find a solution to each target until all the search budget is consumed or all the targets are covered, and (3) a test suite is built combining together all the generated test cases. Therefore, in such techniques, every individual is a test case that evolves to cover a target.

From the ones presented in the literature, we believe that several tools are mature enough to be used in industrial applications, especially for programs written in C language. For instance, Lakhotia et al. implemented AUSTIN, an open-source tool based on the Alternative Variable Method able to deal with pointers [72]. Scalabrino

et al. presented OCELOT, a tool that implements a technique based on the concept of linearly independent path to smartly select the targets and therefore save search budget [110]. Moreover, such a tool is able to automatically generate test cases for the Check³ framework. More recently, Kim et al. introduced CAVM, an extension of a commercial tool able also to handle dynamic data structures [67]. Despite working pretty well for procedural languages, single-target techniques might suffer of some limitations [45]. For instance, in a program under test, some branches might be more difficult to cover or even infeasible. Thus, in this case, a single-target approach would waste a significant amount of budget. Multi-target approaches, discussed in the upcoming paragraph, have been proposed in last years to overcome such limitations.

3.1.2 Multi-Target Approaches

The first example of multi-target technique has been presented by Fraser and Arcuri [45]. They proposed a *whole-suite* approach where the search algorithm evolves the entire suite with the aim to cover all the branches at the same time. In order to achieve such a result, they define a new fitness function that sums the branch distance of all the targets into a cumulative function that express the *goodness* of the entire suite. Such an approach has been implemented in EVOSUITE,⁴ an open-source tool generating JUnit test cases for Java code. Following a similar idea, Panichella et al. proposed MOSA (Many-Objective Sorting Algorithm) [91]. Instead of aggregating multiple objectives into a single value, MOSA reformulates the branch coverage as a many-objective optimization problem. Indeed, in this formulation, a fitness score is a vector of m values, instead of a single aggregate score. In addition, MOSA uses an *archive* to keep track of the best test cases between the many detected by the algorithm. Evaluated on 64 Java classes of large projects, MOSA was able to generate unit test with an average coverage about 84%. Moreover, also such an algorithm is built on top of EVOSUITE. Being available as Maven plugin, such a tool represents an out-of-the-box solution for practitioners that want to automate the process of test case generation.

3.1.3 Limitations and Outlook

The aforementioned approaches only automate the process of generating data able to exercise a part of a system. However, given such input data, a proper test case should be able to check whether the software is behaving as intended, preventing it from potentially incorrect behavior. Such a problem is called the *test oracle problem* [11]. Despite the huge amount of research in testing automation, such a problem

³<https://github.com/libcheck/check>.

⁴<http://www.ev-suite.org>.

still remains less solved. Therefore, without test oracle automation, human effort is needed to determine the correct behavior and inhibits better overall test automation. Moreover, such tools often generate test cases that are hard to understand and difficult to maintain [107]. Despite different approaches that tried to address such a problem, there is still room for improvement.

3.2 *Performance Testing*

Performance testing is a form of nonfunctional testing that deals with the assessment of particular performance counters of a system. A system can range from a piece of software to a deployed application potentially running on multiple computers. Hence, we differentiate between two major types of performance tests: (1) load tests and (2) software microbenchmarks. During load testing, a production-like system is deployed to a dedicated environment, and defined load patterns are executed against that system for a period of time (usually multiple hours). Load testing can be seen as the system/integration testing equivalent for performance. Conversely, microbenchmarking focuses on small fractions of a program (e.g., a function) and evaluates over many executions how performance counters behave for that particular fraction. It is the unit test equivalent for performance. Typical performance counters evaluated are related to time and required resources. Examples are average execution time, throughput, and CPU utilization; and memory consumption, number of allocations, and I/O operations.

In the following, we focus on software microbenchmarking for software-component execution times.

3.2.1 **Problems**

Recent studies on OOS show that microbenchmarking is not as common and popular as unit testing [73, 119]. The decreased popularity is potentially due to multiple factors described in the following. In order to write good performance tests, an in-depth knowledge about compiler/runtime internals and statistics is required [50]. Moreover, execution should be done on an environment dedicated and set up for reliable performance measurements, and tests need to be executed many times (system warmup, >20 measurements) to reduce nondeterministic influences. This results in two constraints: many developers do not have such a dedicated environment but rather use their own machines or unreliable cloud resources, and test-suite execution times rise up to multiple hours or even days [64]. Further, in most programming languages, there is no established standard for writing microbenchmarks, and current tools that support agile process models (i.e., CI servers) do not provide means for continuous performance assessment [73].

3.2.2 Current Solutions

In recent years, software microbenchmarking has gained interest in both academia and industry. To lower the required knowledge for writing microbenchmarks, tool vendors introduced dedicated frameworks that assist in writing good performance tests. OpenJDK introduced from version 7 on the Java Microbenchmarking Harness (JMH).⁵ Newer languages such as Go,⁶ Rust,⁷ and Swift⁸ provide microbenchmarking framework as part of their standard library. On the academic side, Bulej et al. [17] introduce the Stochastic Performance Logic (SPL) that removes the required statistical knowledge from developers for performance test result evaluation. SPL is a declarative way of specifying assertions about a software components performance. One example could be as follows: algorithm A must be faster than algorithm B by a factor of 2. These assertions are transformed into performance tests, and their results are validated with common statistical tests (i.e., hypothesis tests).

Others explored the identification of performance introducing code changes and the reduction of performance test execution time. Jin et al. [64] first study the characteristics of performance bugs and consequently take the insights to compute efficiency rules for performance bug detection. Auxiliary to that, Heger et al. [58] introduce PRCA, an approach that utilizes unit tests and the revision history of a project to find the root cause of a performance problem. Their work bisects the git revision history to find the commit and involved methods that introduced the degradation. Both previously discussed works do not continuously check, but rather check for performance problems ad hoc. Conversely, Huang et al. [61], Alcocer et al. [1], and de Oliveira et al. [35] propose approaches that continuously check for software performance. Huang et al. and Alcocer et al. introduce static approaches to detect potentially performance-risky commits, and based on their assessment flag a commit for benchmark execution or not. Conversely, de Oliveira et al. are the first to introduce a methodology that executes a subset of a performance test suite on every commit. They employ a combination of static analysis, to predict whether a benchmark is potentially able to detect a regression, and historical dynamic benchmark execution data, to predict whether the performance of a benchmark is affected by a commit. Compared to the other approaches, this reduces the benchmark suite to a subset that is of interest and executes a subset on each commit.

3.2.3 Outlook

An unsolved issue so far is the utilization of unstable environments for performance test execution. The premier example of such environments is cloud resources,

⁵<http://openjdk.java.net/projects/code-tools/jmh/>.

⁶<https://golang.org/pkg/testing/>.

⁷<https://doc.rust-lang.org/1.7.0/book/benchmark-tests.html>.

⁸<https://github.com/apple/swift/tree/master/benchmark>.

mostly caused by virtualization and multi-tenancy. Further work in the area of continuously assessing software performance as part of CI needs to be done. We envision a future where performance testing is as common place as unit testing is today, where each build is automatically tested for its performance characteristics, and developers receive quick feedback about these nonfunctional attributes of their software.

4 Continuous Delivery

Continuous Delivery (CD) is an agile software development practice where code changes are released to production in short cycles (i.e., daily or even hourly). A basic CD pipeline is composed by build, deploy, and test stages [123]. This practice is one of the pillars of the agile movement and is widely adopted in both open-source and industry projects by now. The regular invocations of the build-related tools (e.g., static analysis tools) and the corresponding artifacts generated in this process (e.g., build logs) open up new opportunities to better understand the development process and to build tools that support the developers early on.

4.1 *Build Breakage*

In the CD pipeline, a build is typically triggered during the build stage (i.e., Continuous Integration), whenever a code change is pushed in a version control system (e.g., Git). It is being checked out, compiled, tested, and analyzed for code quality measures. The build can potentially fail in any of these phases due to several reasons, for example, syntax errors, failing tests, or violations of coding conventions. Such a failing build delays the release of a new software version. Indeed, developers have to analyze and resolve the problems causing the build failure before being able to perform a new build. In such a scenario, release engineers spend at least 1 h per day to fix broken builds [66] and an organization loses a lot of man-hours because of many build failures occurring during a working day.

Thus, it's crucial to support developers in (1) identifying faster and better the problem causing a build failure and (2) fixing easily those problems.

The first step to meet those two challenges has been a deep understanding of the types of build failures. We performed a large study [124] of 34,182 build failures occurred in OSS and in a large financial organization, namely, ING Nederland. The purpose of this study was to compute a taxonomy of build failures and compare the

frequencies of each category in a closed-source (i.e., ING) and an open-source (i.e., Travis CI⁹) environment.

Through the analysis of the build failures in 349 Java Maven projects from Travis CI and 418 (mostly Java) Maven projects from ING, we derived a Build Failures Catalog including 20 categories. We briefly describe the most important types.

Compilation is the category including builds failed during the compilation of production and test code. A compilation of a code change might not succeed because of language constructs unsupported by the build environment or due to annotations unsupported by the installed Java VM.

Dependencies are another substantial source of build failures. Typical errors in this category are invalid resource configurations or failed downloads due to unavailable artifacts.

Testing failures also break the build. This category is divided in other sub-categories based on the testing activity involved in the failure (e.g., unit testing, integration testing, nonfunctional testing).

Code Analysis enforces code changes to follow predefined code quality criteria. Thus, a build might fail because of non-passed quality gates.

Deployment of artifacts resulting from an introduced code change causes other types of build failures. The deployment environment might be set up incorrectly or the new application version simply doesn't harmonize.

Given this catalog of build failures, we compute and compare the percentages of build failures of different types for both industrial (i.e., ING) and OSS projects. Then we analyzed differences and commonalities. Except for dependencies, we observed a quite different distribution of build failure types in the two domains under analysis.

Specifically, integration testing failures are more frequent in industry than in OSS. Instead, OSS projects exhibit more unit testing failures. Those results suggest that industrial developers are more keen on performing unit tests also before the build (see "Testing" in Fig. 1) and rely on the build server to catch mostly integration issues. In OSS, developers tend to delegate all testing activities to the build stage.

For business-critical projects like the one used by financial organizations, proper nonfunctional testing is crucial. Thus, a separate node is usually used to perform time-consuming testing activities, for example, stress testing and penetration testing. Nevertheless, exclusively industrial developers (at least in ING) started to rely on the build process to spot, whenever possible, nonfunctional issues and specifically load testing failures. It allows them to make the identification of such nonfunctional problems faster and reduce the time and the cost for fixing them.

There are more build failures due to static analysis in industry compared to OSS. Performing a qualitative analysis of both industrial and OSS build failure logs, we observed that most of the OSS projects run static analysis tools directly on the build server, while industry tends to perform static analysis on a different server using

⁹<https://travis-ci.org>.

tool as SonarQube.¹⁰ This choice implies (1) data easy to monitor and query and (2) less overloading of the build server machine.

Finally, we observed a quite low percentage of compilation failures in both domains. This result shows how it's important to compile a code change before building it. This best practice of compiling code before building it makes the identification of the error behind a compilation failure faster and easier (i.e., it's more difficult to spot a compilation error when a code change is already integrated with other changes).

The results of our study suggest the need for supporting developers in maintaining their CD pipeline to make it more efficient, for example, by deciding what to do in private builds on the developer's local machine and what to delegate to build servers, or how to improve the overall detection of the issues by anticipating the execution of nonfunctional tests. We plan to use the taxonomy we built to make the overall process of build failure understanding faster and conceive approaches able to automate the build failure resolution.

4.2 Release Confidence and Velocity

The trend towards highly automated build, test, and deployment processes enables companies delivering their software quickly and efficiently. However, the faster a company moves, the less time is available to perform precautions to minimize the risk of releasing defective changes. Consequently, there exists an inherent trade-off between the risk of lower release quality and time to market. We investigated this trade-off and derived a model [111] based on *release confidence* and the *velocity* of releases during the course of two larger empirical studies [25, 113].

Release Confidence is the amount of confidence gained on the quality gates within a company's development and release process. Those quality gates involve automated (e.g., unit, integration, performance tests) and manual tests (e.g., user acceptance tests) and code reviews.

Release Velocity is the time it takes to assess each single quality gate starting with the developer's commit of a change until the change reaches the production environment, including the time it takes to deploy the newest version.

4.2.1 Model of Release Confidence and Velocity

Our model consists of four categories (*cautious*, *balanced*, *problematic*, *madness*), arranged on a grid from both low to high velocity and confidence. The underlying idea and vision of this model are to serve as a vehicle for self-assessment (i.e., what

¹⁰<https://www.sonarqube.org>.

is my company's category) and provide guidelines on how a company can transition to other categories (e.g., increase velocity while keeping confidence high).

Cautious is the category characterized by high release confidence and low velocity. Companies put a high emphasis on testing, including both a well-maintained set of automated tests to reduce the risk of human-caused errors during manual testing and supplemental manual tests to cover areas hard to test. Code reviews are a common practice complementing the testing phases. Manual approval processes (e.g., domain-specific requirements, company policies) decrease velocity and hence reduce the number of releases.

Problematic is characterized by low confidence in a company's quality gates and low velocity. Typically, this is not a category a company is placed in by choice. Insufficiently maintained test suites (automated and manual), or test suites not covering all aspects of an application, shortages in testing personnel, the absence of code reviews, and unclear roles regarding the quality assurance are characteristics for this category. Velocity is often low as a direct consequence of this, but also due to lack of automation and architectural issues.

Madness is associated with high velocity and low confidence. Release cycles are short; companies make use of early customer feedback and reduced time to market. However, quality assurance plays a minor role, but often by choice, as companies within this category decide that the benefits of sophisticated quality assurance processes are not worth the investments. Consequently, issues might be fixed fast, but the lack of quality gates leads to risky and stressful releases. This category is often appealing for companies with smaller code bases (e.g., startups) as it enables pushing new functionality fast.

Balanced is characterized by high velocity and how confidence and portrays the vision of continuous delivery and deployment [22, 62]. Companies in this category strive for a balance of having sophisticated and highly automated quality assurance processes and code reviews (for specific critical code sections) to maintain confidence on a high level and tool support that allows releasing by the push of a button. Moreover, this category provides a proper basis for post-deployment quality assurance techniques (i.e., continuous experimentation), testing new functionality on a small fraction of the user base first [113].

4.2.2 Transitioning Between Categories

The derived model serves as a starting point discussing the consequences of the categories and allows investigating research gaps not only on how we can better support companies but rather on how to guide (i.e., transition) them to other categories.

Increasing Velocity One topic of raising popularity is containerization, and especially Docker. Docker allows packaging an application with its dependencies into a standardized, self-contained unit that can be used throughout development and to run on any system, being it the development machine, but also the production server.

Its concept of lightweight virtualization speeds up the process of bringing a change into the production system without having to deal with different hardware and software platforms and their dependencies. In a recent study [26] we investigated the Docker ecosystem on GitHub to understand its evolution and identify quality issues. One of the findings is that there is space for improvement when it comes to the size of Docker images. Many projects rely on rather heavyweight OS images as their base image, which somehow defeats the original purpose of lightweight containerization. Therefore, research should aim for providing guidelines and tool support allowing projects to reduce their image sizes and consequently reduce memory consumption when deployed at scale.

Increasing Confidence Recently, the field of continuous experimentation has received increased attention by both academic research and industry (e.g., Fabijan et al. [41], Kohavi et al. [70]). The ability of experimenting with new functionality on a small fraction of the user base enables companies getting early feedback from real-world users while at the same time keeping the risk manageable in case that something goes wrong. Tooling to support experimentation includes our own tool called Bifrost [112], Vamp,¹¹ and Spinnaker.¹² Bifrost and Vamp support the automated, data-driven execution of experiments defined in a domain-specific language. Spinnaker serves as an extension to a CI system allowing to define additional steps for experimentation.

5 Deployment

After completing several iterations of the development cycle introduced in the beginning of this chapter, the updated product will eventually be deployed to its users. While professional software developers are mostly concerned with the quality and state of the current and upcoming releases, it is often desirable to reflect on the long-term evolution of a software project. Managers and project coordinators may be interested in learning how different parts of a project evolve, for example, to reallocate resources and estimate future effort [31, 69, 78, 121]. Modern version control systems present a rich opportunity for understanding the history of a project, but the wealth of information contained within them needs to be managed appropriately.

Two of the main challenges in analyzing the history of a project are the computational time and resource requirements. Running a static analysis tool, for example, to detect bugs or compute various software metrics, can easily take several minutes for a single release. Repeating the analysis for hundreds or even thousands

¹¹<https://vamp.io/>.

¹²<https://www.spinnaker.io/>.

of past releases quickly becomes infeasible. Research has yielded two main avenues of solving this problem: (1) scaling analyses via additional resources, such as clustered computations, and (2) increasing the efficiency of analyses by reducing redundant computation.

A prime example of the former is BOA [40], a server framework that allows analysts to formulate and execute analyses which are executed on a Hadoop cluster. BOA supports analyzing metadata of historical commits (e.g., to learn more about how developers have authored code in the past). It also supports the analysis of Java ASTs to do perform static analyses. For industrial companies, a setup such as this can be useful if there is a large volume of historical code to be analyzed on a regular basis—otherwise, the resources may lay idle.

An alternative to adding more resources is to reduce redundancies during the analyses of past revisions as much as possible. This is done, for example, by LISA, a stand-alone library for running arbitrary analyses on multi-revision graphs [2]. Changes between two subsequent revisions of a software project typically concern only a fraction of the source code, while most of the source code remains identical. LISA exploits this fact by loading source code ASTs and other graphs in such a manner that each node (which might, e.g., represent a Java class) is stored only once for any range of revisions where there have been no changes. While analyzing the graphs, for example, to compute code metrics, computations need only be run once for these entire ranges. This reduces the average time per revision by multiple orders of magnitude compared to the naive approach of analyzing revisions individually. This approach is advantageous when resources cannot be permanently assigned for the analysis of historical data.

No matter which way historical data is computed, the goal is to obtain actionable knowledge on the health of a project. An example for how this can be achieved is Evolizer [48], a library that draws data from both version control systems as well as bug trackers and enables the joined analysis of both resources. Evolizer has, for example, been used to link commits with bug tracking data to automatically determine which parts of the source code are more bug prone, since commits to the same file referencing bugs more frequently are likely to be more fragile. Evolizer has also been used to discover which parts of the source code coevolve and are thus logically coupled.

Current research is addressing the challenges in obtaining linked, historical data efficiently, and it has shown that using this data can provide valuable and useful indicators that can allow software developers to manage the complexity inherent in modern software systems. Exploiting historical data like this allows them to make decisions more confidently and increase the effectiveness of allocated resources, improving the quality and reliability of software in the long term.

6 Summarization Techniques for Code, Change, Testing, Software Reuse, and User Feedback

In the current software industry, developers are involved in a fierce competition to acquire and retain users. Thus, in this competitive market, understanding the factors affecting users' experience and satisfaction and how these factors are related to software quality represents a valuable benefit for developers interested to evolve their software applications [52]. Moreover, with the introduction of Continuous Development and Continuous Integration software practices, it is becoming important for software developers to speed up development activities without hindering the reliability and quality of the produced software [23]. Thus, for “modern software companies it is nowadays, crucial to enact a software development process able (i) to dynamically react to market requirements (i.e., users requests), (ii) delivering at same time high quality and reliable software”. To achieve this high-level goal, developers have to efficiently deal with the huge amount of heterogeneous data [86, 96] they have at their disposal, for example, bug reports [89, 103], source code [14, 56, 75, 81, 115, 116], test cases [44, 94, 123], mailing lists [36, 89], question and answer site (Q&A) discussions [122, 126], user feedback [37, 38, 92], and other kinds of development artifacts [86, 96].

According to its original definition or concept, a *Summarization approach* has the general capability of automatically **extracting** or **abstracting** key content from one or more sources of information [55], thus determining the relevant information in the source being summarized and reducing its content (see Fig. 2). Specifically, the *extraction* capability consists of “selecting original pieces from the source document and concatenating them to yield a shorter text” while the *abstraction* capability is different as it “paraphrases in more general terms what the text is about” [55]. Both the two categories of summaries can be either *indicative*, *informative*, or *critical*:

- **Indicative summary**: it provides a direct link to the required content or relevant sources to users, so that they can read the provided information more depth.
- **Informative summary**: it has the goal to substitute the original source of information, by mainly assembling the relevant content, presenting it in a new, more concise, and structured form.
- **Critical summary (or review)**: it reports or selects the main opinions or statements related to a specific discussed topic; thus, it brings the most relevant feedback, both positives and negatives, about a given subject discussed in the source document.

Given the great potential of such approaches, in recent years, *Summarization techniques* [55, 86] have been explored by SE researchers to conceive approaches and tools that support developers to deal with the management of such a huge amount of heterogeneous data, coming from different sources of information [80, 86, 94, 96].

Fig. 2 High-level view of summarization approaches

In this section, we provide an overview of the summarization techniques explored in the literature for supporting developers during program comprehension, development, maintenance, and testing tasks, by leveraging the abovementioned heterogeneous data. A more detailed and exhaustive literature review on summarization techniques proposed in SE research is reported in recent work [80, 86].

6.1 Source Code Summarization

First attempts on the adoption of summarization techniques to SE problems were of Murphy [84] and Sridhara [115].

The work by Murphy [84] was the first that proposed to generate summaries of source code by analyzing its structural information. In particular, he proposed to use summarization techniques of such techniques for generating abstractive descriptions of its behavior, to automatically document or re-document source code with the generated summaries. Sridhara [115] extended such previous work by Murphy, suggesting the use of predefined *natural language templates*, filled with the main linguistic elements (e.g., verbs, nouns, etc.) [33, 34] composing the signature of methods, to generate the summaries.

On top of such previous work, other researchers used a similar strategy to summarize other kinds of software artifacts at different levels of abstraction: parameters [118], groups of statements [117], Java methods [56, 75, 116], Java classes [81], and services of Java packages [57].

The main limit of such work is that they generate source code descriptions or summaries by only analyzing the static information available in the source code itself. Thus, they are not able to describe the high-level behavior and meaning of the described software artifact, something that developers often report in various communication means, such as mailing list [10, 90], issue trackers [89, 90, 104], IRC chat log [90], and other developers communication means. For this reason, recent work proposed to generate source code documentation by mining text data from other sources of information, alternative to source code: question and answer site (Q&A) discussions [122, 126], bug reports [89, 104], e-mails [89], and forum posts.

6.2 Task-Driven Software Summarization

A limit of approaches proposed for *Code Summarization* regards the way in which SE researchers evaluated their usefulness. Indeed, Binkley et al. [14] and Jones et al. [114] highlighted that the evaluation of summarization techniques for SE should not be done by using simple metrics, answering the general question “is this a good summary?”. Thus, they proposed the concept of *Task-Driven Software Summarization* where summarization techniques for SE should be evaluated “through the lens of a particular task” (e.g., during bug fixing or testing tasks). Stemming from the observations made by Binkley et al. [14] and Jones et al. [114], recent work proposed approaches for automating particular software maintenance and testing tasks [80, 86], evaluating their practical usefulness in their specific utilization context.

6.2.1 Code Change Summarization

Code change summarization approaches have the goal of augmenting the context provided by differencing tools, generating natural language descriptions of changes occurred at different types of software artifacts [80, 86].

An example of such differencing tools is the *Semantic Diff* tool introduced by Jackson and Ladd [71] which detects differences between two versions of a procedure and uses program analysis techniques to summarize the semantic differences. More recent and modern examples of differencing tools are *DeltaDoc* [18] and *Commit 2.0* [30] that (1) describe source code modifications using symbolic execution with summarization techniques and (2) augment commit logs with a visual context of the changes, respectively. However, most of code change summarization approaches proposed in the literature [28, 82] are based on a specific,

well-known differencing tool called *Change Distiller*, implemented by Fluri et al. [43], which extracts fine-grained source code changes based on a specialized tree differencing algorithm. Thus, Change Distiller generates a list of classified changes based on the operation type performed by the developers in the analyzed commit, that is, *insertion*, *deletion*, or *modification*. On top of such change types information, it generates the corresponding customized abstract syntax tree.

Hence, information coming from Change Distiller tools have been exploited, for example, by researchers to automatically generate high-quality (1) commit messages [28] and (2) release notes [82]. Specifically, the approach for generating commit messages, called Change Scribe, conceived by Cortes-Coy et al. [28] takes as input two consecutive versions of a Java project; then it (1) uses Change Distiller to extract the source code changes occurred between the two versions of the project (changes related to addition, removal, or modification types); (2) detects responsibilities of methods within each Java class using the concept of method stereotypes; (3) characterizes the change set using commit stereotypes; (4) estimates the impact set for the changes in the commit; (5) performs the selection of the content considering the threshold values defined by the developer; and (6) finally generates the change descriptions for the analyzed commit. The tool proposed by Moreno et al. [82], called ARENA, extended the work by Cortes-Coy et al. [28], extracting the information about the change types from two different sources of information, namely, the versioning system and the source archives of the releases to be compared. ARENA and Change Scribe achieved high accuracy in generating high-quality commit messages and release notes. Indeed, in some cases, according to the involved study participants, they were often preferred to the one written by the original developers.

6.2.2 Summarization Techniques for Testing and Code Reuse

As discussed previously, most of previous works on source code and code change summarization have been evaluated by simply surveying developers about the general quality of the provided summaries [14, 56, 75, 81, 115, 116]. However, recent work proposed summarization techniques to support developers during bug fixing and/or testing tasks [65, 94], demonstrating their practical usefulness in performing such tasks.

Specifically, since *Waterfall* up until *Agile*, Software Testing has been playing an essential role in any software development methodology to detect the defects of software products in different target environments. However, developers perceived testing as a time-consuming task because it requires a quarter of their working time engineering tests [13] and up to 50% of the overall project effort [15]. In this scenario, automated test generation tools [46, 47] in software development pipelines could potentially reduce the time spent by developers in writing test cases. The main advantages of such tools include the generation of tests achieving higher code coverage when compared to the coverage obtained through manual testing [47] and to find violations of undeclared exceptions [46].

Despite such undisputed advantages, nowadays, automated test generation tools are still not used in practice. The main reason is that the generated tests are too hard to understand and difficult to maintain [29, 94]. As a consequence, generated tests do not improve the ability of developers to detect faults when compared to manual testing [20, 47]. Thus, to foster the adoption of automated testing tools, Panichella et al. [94] presented *TestDescriber*, a tool that summarizes both automatically generated or manually written JUnit test cases. Specifically, taking as input the JUnit test and the corresponding class under test (CUT), TestDescriber (i) runs each test method tracing statements and branches exercised in the CUT and (iii) augments the JUnit test with summaries that, at different abstraction levels, provides a dynamic view of the CUT.

Having TestDescriber at their disposal, Panichella et al. [94] performed an empirical study [94] to investigate the usefulness of the proposed tool when used in a concrete scenario of use: a Java class has been modified or developed and must be tested using generated test cases with the purpose of identifying and fixing eventual bugs affecting the production code. Thus, the goal of the study was to determine the impact of the generated test summaries on the number of bugs actually fixed by developers when assisted by automated test generation tools. The results of our study show that participants without TestDescriber summaries were able to remove only 40% of bugs present on the considered classes. Instead, when relying on test case with summaries, TestDescriber improved the bug fixing performance of the participants from 50% up to 100%. The results of the Wilcoxon test highlighted that the result was statistically significant (with p -values always <0.05).

Thus, differently from most work in the literature, the work by Panichella et al. [94] is the first that deals with the defect or bug detection. Indeed, it is important to mention that other recent work in the literature proposed the use of NLP templates and/or summarization techniques mostly to document undocumented part of source code, without addressing the problem of bug detection [80, 86]. Following this line of research, Zhou et al. [130] proposed an approach based on NLP templates, able to detect *API defects* in Java Libraries. Specifically, application programming interfaces (APIs) represent the most adopted tools for developers to build complex software systems nowadays. However, several studies have revealed that also major API providers tend to have an incomplete or inconsistent API documentation. This severely hampers the API comprehension and the quality of software built on it. Thus, Zhou et al. [130] proposed DRONE, a framework to automatically detect and repair defects affecting API documents by leveraging techniques from program analysis, natural language processing, and constraint solving. The research evaluation involving part of well-documented JDK 1.8 APIs has shown that DRONE is able to detect API defects with an average F-measure of 79.9%, 71.7%, and 81.4%, respectively, demonstrating its usefulness.

6.3 Summarization of Textual User Feedback

In the current software industry, developers are involved in a fierce competition to acquire and retain users. Thus, in this competitive market, understanding the factors affecting users' experience and satisfaction and how these factors are related to software quality represents a valuable benefit for developers interested to evolve their software applications. In this context, app stores, such as Google Play or the Apple Store, allow users to provide feedback on apps by posting review comments and giving star ratings. The experience an end user has with apps reported in user reviews is a key concern when creating and maintaining any successful application. For this reason, mobile developers are interested to exploiting opinions and/or feedback of end users during the evolution of their software [24, 125].

As discussed previously, automatically generated summaries can be either *indicative*, *informative*, or *critical* [55]. Specifically, a **critical summary (or review)** reports or selects the main opinions or statements related to a specific discussed topic; thus, it brings the most relevant feedback, both positives and negatives, about a given subject discussed in the source document. Thus, the peculiarity of critical summaries has pushed researchers to conceive tools for automatically extracting user feedback from user review, relevant for software evolution [24, 27, 53, 54, 63, 74, 92, 93]. For instance, Chen et al. presented a computational framework which automatically groups, prioritizes, and visualizes informative reviews [24]. However, most of proposed tools only perform a simple classification of user reviews according to specific topics [53, 54, 63, 88, 92, 92], without reducing the amount of reviews developers have to deal with, which is very large for popular apps.

A more recent work by Ciurumelea et al. [27] proposed an approach that classifies reviews according to more fine-grained topics addressed by users in app reviews, which uses machine learning to classify reviews according to such topics. In addition, to deal with such an amount of user review data, Di Sorbo et al. [37, 38] proposed an approach called SURF, which at the same time (1) determines the specific topic discussed in the review (e.g., UI improvements, security issues, etc.), (2) identifies the maintenance task to perform for addressing the request stated in the review (e.g., bug fixing, feature enhancement, etc.), and (3) presents such information to developers as an actionable *condensed, interactive, and structured agenda of recommended software changes*. The approach relies on a conceptual model of the user requests reported by app reviews; then it uses sophisticated summarization approaches, based on machine learning and NLP techniques, for summarizing thousands of reviews. Di Sorbo et al. [37, 38] performed an end-to-end evaluation of the proposed approach on user reviews of 17 mobile apps involving 23 developers. Results demonstrate high accuracy of SURF in summarizing reviews containing feedback for planning future software changes, substantially reducing the time and effort required for manually analyzing user review content.

6.4 Future Research

Recent research in SE observed an increasing adoption of summarization techniques for accomplishing simple or more complex development, maintenance, and testing tasks. However, their adoption in industrial contexts requires substantial novel and advanced research to make them applicable in any industrial or open-source organizations. Thus, future research in the field is devoted to fill the existing gap between industrial needs and current provided research prototypes:

- *Summarization of heterogeneous data*: current summarization approaches are mostly conceived for analyzing one or two sources of information. However, when performing development, maintenance, and testing tasks, developers access various types of heterogeneous data. Thus, future summarization techniques should be designed with advanced mechanisms able to distill, in a simultaneous manner, the relevant knowledge present in different sources of information, presenting it in a unified manner, depending on the specific task the developer is performing.
- *Scalability and integration in the CD/CI process*: current summarization approaches are able to proficiently distill relevant information from various kinds of software artifacts. However, they are usually computationally expensive and thus not applicable in real working contexts. Moreover, most of such tools are difficult to integrate in the current continuous delivery software development process. Hence, future research should be devoted on designing tools able to analyze, with substantial low computational cost, the huge amount of available heterogeneous data, integrating the summarized information in the various development phases composing the CD pipeline applied in a software organization.
- *Visualization of software summaries*: most of the generated software summaries are presented as a set of textual fragments that share similar concepts. Thus, part of the future research related to the application of source code and code change summarization needs to be devoted to the definition of proper visualization metaphors that actually present the information provided by the generated summaries in a more structured manner.

7 Summary

Today's software development is about data and processes. Environments and tools are cornerstones for any successful project. At the same time, we need to adopt new techniques from research to deal with this sheer amount of information that ranges from requirements to code, to tests, and to deployment and experimentation.

We discussed several techniques and their potential for these data-driven decisions and actions that are ready for use in practice. The techniques that we presented originate mostly from our own research and provide new solutions

to recommender systems, automated test case generation, performance testing, continuous integration and continuous deployment, evolution analysis feedbacks, and summarization techniques for code or tests.

The next step is to adopt and integrate these and other techniques in the daily development activities and engineering processes. We showed some of the great potential of these technologies in dealing with the large amount of data that is available in any process step of software development.

New development environments are rising that support live development (in the cloud). Such tools require data that is beyond static code or test data; we need to build active feedback loops into the programming environments that employ proper data analytics customized for each development step. Tools such as live programming or programming in the cloud will then be eased and be put on a much more stable basis. With the increasing speed of software delivery to customers, these links between the developer and the customer need to become seamless and active, building on the data aggregated, accumulated, and analyzed for recommendations, summarizations, or testing in various dimensions.

References

1. Alcocer, J.P.S., Bergel, A., Valente, M.T.: Learning from source code history to identify performance failures. In: Proceedings of the 7th ACM/ SPEC International Conference on Performance Engineering (ICPE), pp. 37–48 (2016)
2. Alexandru, C.V., Panichella, S., Gall, H.: Reducing redundancies in multi-revision code analysis. In: IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER), Klagenfurt (2017)
3. Ali, S., Briand, L.C., Hemmati, H., Panesar-Walawege, R.K.: A systematic review of the application and empirical investigation of search-based test case generation. *IEEE Trans. Softw. Eng.* **36**(6), 742–762 (2010)
4. Allamanis, M., Sutton, C.: Mining source code repositories at massive scale using language modeling. In: Proceedings of the 10th Working Conference on Mining Software Repositories, MSR '13, pp. 207–216. IEEE Press, Piscataway (2013)
5. Allamanis, M., Barr, E.T., Bird, C., Sutton, C.: Learning natural coding conventions. In: Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2014, pp. 281–293. ACM, New York (2014)
6. Allamanis, M., Barr, E.T., Bird, C., Sutton, C.: Suggesting accurate method and class names. In: Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015, pp. 38–49. ACM, New York (2015)
7. Allamanis, M., Barr, E.T., Devanbu, P., Sutton, C.: A Survey of Machine Learning for Big Code and Naturalness. *arXiv e-prints*, September 2017
8. Amann, S., Proksch, S., Nadi, S.: FeedBaG: an interaction tracker for visual studio. In: International Conference on Program Comprehension. IEEE, Piscataway (2016)
9. Amann, S., Proksch, S., Nadi, S., Mezini, M.: A study of visual studio usage in practice. In: International Conference on Software Analysis, Evolution, and Reengineering. IEEE, Piscataway (2016)
10. Bacchelli, A., Sasso, T.D., D'Ambros, M., Lanza, M.: Content classification of development emails. In: 34th International Conference on Software Engineering, ICSE 2012, June 2–9, Zurich, pp. 375–385 (2012)

11. Barr, E.T., Harman, M., McMinn, P., Shahbaz, M., Yoo, S.: The oracle problem in software testing: a survey. *IEEE Trans. Softw. Eng.* **41**(5), 507–525 (2015)
12. Beizer, B.: *Software Testing Techniques*, 2nd edn. Van Nostrand Reinhold Co., New York (1990)
13. Beller, M., Gousios, G., Panichella, A., Zaidman, A.: When, how, and why developers (do not) test in their IDEs. In: *Proceedings of the 10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE)*. ACM, New York (2015)
14. Binkley, D., Lawrie, D., Hill, E., Burge, J., Harris, I., Hebig, R., Keszocze, O., Reed, K., Slankas, J.: Task-driven software summarization. In: *Proceedings of the International Conference on Software Maintenance (ICSM)*, pp. 432–435. IEEE, Piscataway (2013)
15. Brooks, F.P.Jr.: *The Mythical Man-Month*. Addison-Wesley, Reading (1975)
16. Bruch, M., Monperrus, M., Mezini, M.: Learning from examples to improve code completion systems. In: *Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering*, pp. 213–222. ACM, New York (2009)
17. Bulej, L., Bureš, T., Horký, V., Kotrč, J., Marek, L., Trojánek, T., Tůma, P.: Unit testing performance with stochastic performance logic. *Autom. Softw. Eng.* **24**(1), 139–187 (2017)
18. Buse R.P.L., Weimer, W.R.: Automatically documenting program changes. In: *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering, ASE '10*, pp. 33–42. ACM, New York (2010)
19. Campbell, J.C., Hindle, A., Amaral, J.N.: Syntax errors just aren't natural: improving error reporting with language models. In: *Proceedings of the 11th Working Conference on Mining Software Repositories, MSR 2014*, pp. 252–261. ACM, New York (2014)
20. Ceccato, M., Marchetto, A., Mariani, L., Nguyen, C.D., Tonella, P.: Do automatically generated test cases make debugging easier? An experimental assessment of debugging effectiveness and efficiency. *ACM Trans. Softw. Eng. Methodol.* **25**(1), 5:1–5:38 (2015)
21. Chang, K.H., Cross II, J.H., Carlisle, W.H., Liao, S.-S.: A performance evaluation of heuristics-based test case generation methods for software branch coverage. *Int. J. Softw. Eng. Knowl. Eng.* **6**(04), 585–608 (1996)
22. Chen, L.: Continuous delivery: huge benefits, but challenges too. *Softw. IEEE* **32**(2), 50–54 (2015)
23. Chen, L.: Continuous delivery: overcoming adoption challenges. *J. Syst. Softw.* **128**, 72–86 (2017)
24. Chen, N., Lin, J., Hoi, S.C.H., Xiao, X., Zhang, B.: Ar-miner: mining informative reviews for developers from mobile app marketplace. In: *Proceedings of the 36th International Conference on Software Engineering, ICSE 2014*, pp. 767–778. ACM, New York (2014)
25. Cito, J., Leitner, P., Fritz, T., Gall, H.C.: The making of cloud applications: an empirical study on software development for the cloud. In: *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering (ESEC/FSE)*, pp. 393–403. ACM, New York (2015)
26. Cito, J., Schermann, G., Wittern, J.E., Leitner, P., Zumberi, S., Gall, H.C.: An empirical analysis of the docker container ecosystem on github. In: *Proceedings of the 14th International Conference on Mining Software Repositories, MSR '17*, pp. 323–333. IEEE Press, Piscataway (2017)
27. Ciurumelea, A., Schauffelbühl, A., Panichella, S., Gall, H.: Analyzing reviews and code of mobile apps for better release planning. In: *2017 IEEE 24th IEEE International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, pp. 91–102 (2017)
28. Cortes-Coy, L.F., Vázquez, M.L., Aponte, J., Poshyvanyk, D.: On automatically generating commit messages via summarization of source code changes. In: *Proceedings of the International Working Conference on Source Code Analysis and Manipulation (SCAM)*, pp. 275–284. IEEE, Piscataway (2014)
29. Daka, E., Campos, J., Fraser, G., Dorn, J., Weimer, W.: Modeling readability to improve unit tests. In: *Proceedings of the 10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE)*. ACM, New York (2015)

30. D'Ambros, M., Lanza, M., Robbes, R.: Commit 2.0. In: Proceedings of the 1st Workshop on Web 2.0 for Software Engineering, Web2SE '10, pp. 14–19. ACM, New York (2010)
31. D'Ambros, M., Lanza, M., Robbes, R.: Evaluating defect prediction approaches: a benchmark and an extensive comparison. *Empir. Softw. Eng.* **17**(4–5), 531–577 (2012)
32. Damevski, K., Shepherd, D., Schneider, J., Pollock, L.: Mining sequences of developer interactions in visual studio for usage smells. *IEEE Trans. Softw. Eng.* **43**(4), 359–371 (2016)
33. De Lucia, A., Di Penta, M., Oliveto, R., Panichella, A., Panichella, S.: Using IR methods for labeling source code artifacts: is it worthwhile? In: IEEE 20th International Conference on Program Comprehension, ICPC 2012, Passau, June 11–13, 2012, pp. 193–202 (2012)
34. De Lucia, A., Di Penta, M., Oliveto, R., Panichella, A., Panichella, S.: Labeling source code with information retrieval methods: an empirical study. *Empir. Softw. Eng.* **19**(5), 1383–1420 (2014)
35. de Oliveira, A.B., Fischmeister, S., Diwan, A., Hauswirth, M., Sweeney, P.: Perphhecy: performance regression test selection made simple but effective. In: Proceedings of the 10th IEEE International Conference on Software Testing, Verification and Validation (ICST), Tokyo (2017)
36. Di Sorbo, A., Panichella, S., Visaggio, C.A., Di Penta, M., Canfora, G., Gall, H.C.: Development emails content analyzer: intention mining in developer discussions. In: 2015, 30th IEEE/ACM International Conference on Automated Software Engineering (ASE), pp. 12–23. IEEE, Washington (2015)
37. Di Sorbo, A., Panichella, S., Alexandru, C., Shimagaki, J., Visaggio, C.A., Canfora, G., Gall, H.C.: What would users change in my app? summarizing app reviews for recommending software changes. In: 2016 ACM SIGSOFT International Symposium on the Foundations of Software Engineering (FSE), pp. 499–510 (2016)
38. Di Sorbo, A., Panichella, S., Alexandru, C.V., Visaggio, C.A., Canfora, G.: Surf: summarizer of user reviews feedback. In: Proceedings of the 39th International Conference on Software Engineering Companion, pp. 55–58. IEEE Press, Piscataway (2017)
39. Dias, M., Cassou, D., Ducasse, S.: Representing code history with development environment events. In: International Workshop on Smalltalk Technologies (2013)
40. Dyer, R.: Bringing ultra-large-scale software repository mining to the masses with Boa. PhD thesis, Ames (2013). AAI3610634
41. Fabijan, A., Dmitriev, P., Olsson, H.H., Bosch, J.: The evolution of continuous experimentation in software product development. In: International Conference on Software Engineering, ICSE, Buenos Aires (2017)
42. Ferguson, R., Korel, B.: The chaining approach for software test data generation. *ACM Trans. Softw. Eng. Methodol.* **5**(1), 63–86 (1996)
43. Fluri, B., Wuersch, M., Pinzger, M., Gall, H.: Change distilling: tree differencing for fine-grained source code change extraction. *IEEE Trans. Softw. Eng.* **33**(11), 725–743 (2007)
44. Fraser, G., Arcuri, A.: Evosuite: automatic test suite generation for object-oriented software. In: Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering, ESEC/FSE '11, pp. 416–419. ACM, New York (2011)
45. Fraser, G., Arcuri, A.: Whole test suite generation. *IEEE Trans. Softw. Eng.* **39**(2), 276–291 (2013)
46. Fraser, G., Arcuri, A.: 1600 faults in 100 projects: automatically finding faults while achieving high coverage with evosuite. *Empir. Softw. Eng.* **20**(3), 611–639 (2015)
47. Fraser, G., Staats, M., McMinn, P., Arcuri, A., Padberg, F.: Does automated white-box test generation really help software testers? In: Proceedings of the International Symposium on Software Testing and Analysis (ISSTA), pp. 291–301. ACM, New York (2013)
48. Gall, H.C., Fluri, B., Pinzger, M.: Change analysis with evolizer and changedistiller. *Software, IEEE* **26**(1), 26–33 (2009)
49. Gallagher, M.J., Lakshmi Narasimhan, V.: Adtest: a test data generation suite for ada software systems. *IEEE Trans. Softw. Eng.* **23**(8), 473–484 (1997)

50. Georges, A., Buytaert, D., Eeckhout, L.: Statistically rigorous java performance evaluation. In: *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications, OOPSLA '07*, pp. 57–76. ACM, New York (2007)
51. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization and Machine Learning*, 1st edn. Addison-Wesley Longman Publishing Co., Inc., Boston (1989)
52. Grano, G., Di Sorbo, A., Mercaldo, F., Aaron Visaggio, C., Canfora, G., Panichella, S.: Android apps and user feedback: a dataset for software evolution and quality improvement. In: *Proceedings of the 2nd ACM SIGSOFT International Workshop on App Market Analytics, WAMA@ESEC/SIGSOFT FSE 2017*, Paderborn, September 5, 2017, pp. 8–11 (2017)
53. Guzman, E., Maalej, W.: How do users like this feature? A fine grained sentiment analysis of app reviews. In: *2014 IEEE 22nd International Requirements Engineering Conference (RE)*, pp. 153–162 (2014)
54. Ha, E., Wagner, D.: Do android users write about electric sheep? Examining consumer reviews in google play. In: *Consumer Communications and Networking Conference (CCNC)*, 2013 IEEE, pp. 149–157 (2013)
55. Hahn, U., Mani, I.: The challenges of automatic summarization. *Computer* **33**(11), 29–36 (2000)
56. Haiduc, S., Aponte, J., Moreno, L., Marcus, A.: On the use of automated text summarization techniques for summarizing source code. In: *Proceedings of the International Working Conference on Reverse Engineering (WCRE)*, pp. 35–44. IEEE, New York (2010)
57. Hammad, M., Abuljadayel, A., Khalaf, M.: Automatic summarising: the state of the art. *Lect. Notes Softw. Eng.* **4**(2), 129–132 (2016)
58. Heger, C., Happe, J., Farahbod, R.: Automated root cause isolation of performance regressions during software development. In: *Proceedings of the 4th ACM/SPEC International Conference on Performance Engineering, ICPE '13*, pp. 27–38. ACM, New York (2013)
59. Hellendoorn, V.J., Devanbu, P.T., Bacchelli, A.: Will they like this? Evaluating code contributions with language models. In: *Proceedings of the 12th Working Conference on Mining Software Repositories, MSR '15*, pp. 157–167. IEEE Press, Piscataway (2015)
60. Hindle, A., Barr, E.T., Su, Z., Gabel, M., Devanbu, P.: On the naturalness of software. In: *Proceedings of the 34th International Conference on Software Engineering, ICSE '12*, pp. 837–847. IEEE Press, Piscataway (2012)
61. Huang, P., Ma, X., Shen, D., Zhou, Y.: Performance regression testing target prioritization via performance risk analysis. In: *Proceedings of the 36th International Conference on Software Engineering, ICSE 2014*, pp. 60–71. ACM, New York (2014)
62. Humble, J., Farley, D.: *Continuous Delivery: Reliable Software Releases Through Build, Test, and Deployment Automation*. Addison-Wesley Professional, Reading (2010)
63. Iacob, C., Harrison, R.: Retrieving and analyzing mobile apps feature requests from online reviews. In: *Proceedings of the 10th Working Conference on Mining Software Repositories, MSR'13*, pp. 41–44. IEEE Press, Piscataway (2013)
64. Jin, G., Song, L., Shi, X., Scherpelz, J., Lu, S.: Understanding and detecting real-world performance bugs. In: *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12*, pp. 77–88. ACM, New York (2012)
65. Kamimura, M., Murphy, G.C.: Towards generating human-oriented summaries of unit test cases. In: *Proceedings of the International Conference on Program Comprehension (ICPC)*, May, pp. 215–218. IEEE, Piscataway (2013)
66. Kerzazi, N., Khomh, F., Adams, B.: Why do automated builds break? An empirical study. In: *30th IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pp. 41–50. IEEE, Piscataway (2014)
67. Kim, J., You, B., Kwon, M., McMinn, P., Yoo, S.: Evaluating CAVM: a new search-based test data generation tool for C. In: *International Symposium on Search-Based Software Engineering (SSBSE 2017)* (2017)
68. Ko, A.J., Myers, B.A., Aung, H.H.: Six learning barriers in end-user programming systems. In: *2004 IEEE Symposium on Visual Languages and Human Centric Computing*, pp. 199–206. IEEE, Washington (2004)

69. Kocaguneli, E., Menzies, T., Keung, J.W.: On the value of ensemble effort estimation. *IEEE Trans. Softw. Eng.* **38**(6), 1403–1416 (2012)
70. Kohavi, R., Deng, A., Frasca, B., Walker, T., Xu, Y., Pohlmann, N.: Online controlled experiments at large scale. In: *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1168–1176. ACM, New York (2013)
71. Lahiri, S.K., Hawblitzel, C., Kawaguchi, M., Rebêlo, H.: SYMDIFF: A Language-Agnostic Semantic Diff Tool for Imperative Programs, pp. 712–717. Springer, Berlin (2012)
72. Lakhotia, K., Harman, M., Gross, H. (2013) Austin: an open source tool for search based software testing of C programs. *Inf. Softw. Technol.* **55**(1), 112–125 (2013)
73. Leitner, P., Bezemer, C.-P.: An exploratory study of the state of practice of performance testing in java-based open source projects. In: *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering, ICPE '17*, pp. 373–384. ACM, New York (2017)
74. Maalej, W., Nabil, H.: Bug report, feature request, or simply praise? On automatically classifying app reviews. In: *2015 IEEE 23rd International Requirements Engineering Conference (RE)*, August, pp. 116–125 (2015)
75. McBurney, P.W., McMillan, C.: Automatic documentation generation via source code summarization of method context. In: *Proceedings of the International Conference on Program Comprehension (ICPC)*, pp. 279–290. ACM, New York (2014)
76. McCandless, M., Hatcher, E., Gospodnetic, O.: *Lucene in Action: Covers Apache Lucene 3.0*. Manning Publications Co., Greenwich (2010)
77. McMin, P.: Search-based software testing: past, present and future. In: *Proceedings of the 2011 IEEE Fourth International Conference on Software Testing, Verification and Validation Workshops, ICSTW '11*, pp. 153–163. IEEE Computer Society, Washington (2011)
78. Mende, T., Koschke, R.: Revisiting the evaluation of defect prediction models. In: *Proceedings of the 5th International Conference on Predictor Models in Software Engineering, PROMISE '09*, pp. 7:1–7:10. ACM, New York (2009)
79. Minelli, R., Mocci, A., Robbes, R., Lanza, M.: Taming the ide with fine-grained interaction data. In: *International Conference on Program Comprehension* (2016)
80. Moreno, L., Marcus, A.: Automatic software summarization: the state of the art. In: *Proceedings of the 39th International Conference on Software Engineering, ICSE 2017, Buenos Aires, May 20–28, 2017—Companion Volume*, pp. 511–512 (2017)
81. Moreno, L., Aponte, J., Sridhara, G., Marcus, A., Pollock, L., Vijay-Shanker, K.: Automatic generation of natural language summaries for java classes. In: *Proceedings of the International Conference on Program Comprehension (ICPC)*, May, pp. 23–32. IEEE, Piscataway (2013)
82. Moreno, L., Bavota, G., Di Penta, M., Oliveto, R., Marcus, A., Canfora, G.: Automatic generation of release notes. In: *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, (FSE-22)*, Hong Kong, November 16–22, 2014, pp. 484–495 (2014)
83. Moritz, E., Linares-Vásquez, M., Poshvanyk, D., Grechanik, M., McMillan, C., Gethers, M.: Export: detecting and visualizing API usages in large source code repositories. In: *Proceedings of the 28th IEEE/ACM International Conference on Automated Software Engineering*, pp. 646–651. IEEE Press, Piscataway (2013)
84. Murphy, G.C.: *Lightweight structural summarization as an aid to software evolution*. PhD thesis (1996). AAI9704521
85. Murphy, G.C., Kersten, M., Findlater, L.: How are java software developers using the eclipse IDE? *IEEE Softw.* **23**(4), 76–83 (2006)
86. Nazar, N., Hu, Y., Jiang, H.: Summarizing software artifacts: a literature review. *J. Comput. Sci. Technol.* **31**(5), 883–909 (2016)
87. Negara, S., Vakilian, M., Chen, N., Johnson, R.E., Dig, D.: Is it dangerous to use version control histories to study source code evolution? In: *European Conference on Object-Oriented Programming*, pp. 79–103. Springer, Heidelberg (2012)

88. Palomba, F., Salza, P., Ciurumelea, A., Panichella, S., Gall, H.C., Ferrucci, F., De Lucia, A.: Recommending and localizing change requests for mobile apps based on user reviews. In: *Proceedings of the 39th International Conference on Software Engineering, ICSE 2017*, Buenos Aires, May 20–28, pp. 106–117 (2017)
89. Panichella, S., Aponte, J., Di Penta, M., Marcus, A., Canfora, G.: Mining source code descriptions from developer communications. In: *Proceedings of the International Conference on Program Comprehension, ICPC*, pp. 63–72. IEEE, Los Alamitos (2012)
90. Panichella, S., Bavota, G., Di Penta, M., Canfora, G., Antoniol, G.: How developers' collaborations identified from different sources tell us about code changes. In: *30th IEEE International Conference on Software Maintenance and Evolution*, Victoria, September 29–October 3, pp. 251–260 (2014)
91. Panichella, A., Kifetew, F.M., Tonella, P.: Reformulating branch coverage as a many-objective optimization problem. In: *ICST*, pp. 1–10. IEEE Computer Society, Washington (2015)
92. Panichella, S., Di Sorbo, A., Guzman, E., Visaggio, C.A., Canfora, G., Gall, H.C.: How can I improve my app? Classifying user reviews for software maintenance and evolution. In: *2015 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pp. 281–290 (2015)
93. Panichella, S., Di Sorbo, A., Guzman, E., Visaggio, C.A., Canfora, G., Gall, G., Gall, H.C.: Ardoc: app reviews development oriented classifier. In: *2016 ACM SIGSOFT International Symposium on the Foundations of Software Engineering (FSE)*, pp. 1023–1027 (2016)
94. Panichella, S., Panichella, A., Beller, M., Zaidman, A., Gall, H.C.: The impact of test case summaries on bug fixing performance: an empirical investigation. In: *Proceedings of the 38th International Conference on Software Engineering, ICSE '16*, pp. 547–558. ACM, New York (2016)
95. Ponzanelli, L., Bavota, G., Di Penta, M., Oliveto, R., Lanza, M.: Prompter: a self-confident recommender system. In: *2014 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pp. 577–580. IEEE, Victoria (2014)
96. Ponzanelli, L., Mocci, A., Lanza, M.: Summarizing complex development artifacts by mining heterogeneous data. In: *Proceedings of the 12th Working Conference on Mining Software Repositories, MSR '15*, pp. 401–405. IEEE Press, Piscataway (2015)
97. Proksch, S., Bauer, V., Murphy, G.C.: How to build a recommendation system for software engineering. In: *Software Engineering*. Springer, Berlin (2015)
98. Proksch, S., Lerch, J., Mezini, M.: Intelligent code completion with Bayesian networks. *Trans. Softw. Eng. Methodol.* **25**(1), 3 (2015)
99. Proksch, S., Amann, S., Nadi, S., Mezini, M.: Evaluating the evaluations of code recommender systems: a reality check. In: *International Conference on Automated Software Engineering*. ACM, New York (2016)
100. Proksch, S., Amann, S., Nadi, S.: Enriched event streams: a general dataset for empirical studies on in-IDE activities of software developers. In: *International Conference on Mining Software Repositories (accepted Mining Challenge)* (2017)
101. Proksch, S., Nadi, S., Amann, S., Mezini, M.: Enriching in-IDE process information with fine-grained source code history. In: *International Conference on Software Analysis, Evolution, and Reengineering* (2017)
102. Radevski, S., Hata, H., Matsumoto, K.: Towards building api usage example metrics. In: *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, vol. 1, pp. 619–623. IEEE, Piscataway (2016)
103. Rastkar, S., Murphy, G.C., Murray, G.: Summarizing software artifacts: a case study of bug reports. In: *Proceedings of the 32Nd ACM/IEEE International Conference on Software Engineering - volume 1, ICSE '10*, pp. 505–514 (2010)
104. Rastkar, S., Murphy, G.C., Murray, G.: Automatic summarization of bug reports. *IEEE Trans. Softw. Eng.* **40**(4), 366–380 (2014)
105. Ray, B., Hellendoorn, V., Godhane, S., Tu, Z., Bacchelli, A., Devanbu, P.: On the “naturalness” of buggy code. In: *Proceedings of the 38th International Conference on Software Engineering, ICSE '16*, pp. 428–439. ACM, New York (2016)

106. Robillard, M.P.: What makes APIs hard to learn? Answers from developers. *IEEE Softw.* **26**(6), 27–39 (2009)
107. Rojas, J.M., Fraser, G., Arcuri, A.: Automated unit test generation during software development: a controlled experiment and think-aloud observations. In: *Proceedings of the 2015 International Symposium on Software Testing and Analysis, ISSTA 2015*, pp. 338–349. ACM, New York (2015)
108. Saied, M.A., Benomar, O., Abdeen, H., Sahraoui, H.: Mining multi-level api usage patterns. In: *2015 IEEE 22nd International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pp. 23–32. IEEE, Piscataway (2015)
109. Saied, M.A., Benomar, O., Abdeen, H., Sahraoui, H.: Mining multi-level api usage patterns. In: *2015 IEEE 22nd International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pp. 23–32. IEEE, Piscataway (2015)
110. Scalabrino, S., Grano, G., Di Nucci, D., Oliveto, R., De Lucia, A.: Search-based testing of procedural programs: iterative single-target or multi-target approach? In: *Search Based Software Engineering*, October, pp. 64–79. Springer, Cham (2016)
111. Schermann, G., Cito, J., Leitner, P., Zdon, U., Gall, H.C.: Towards quality gates in continuous delivery and deployment. In: *2016 IEEE 24th International Conference on Program Comprehension (ICPC)*, pp. 1–4. IEEE, Piscataway (2016)
112. Schermann, G., Schöni, D., Leitner, P., Gall, H.C.: Bifrost: supporting continuous deployment with automated enactment of multi-phase live testing strategies. In: *Proceedings of the 17th International Middleware Conference*, pp. 12:1–12:14. ACM, New York (2016)
113. Schermann, G., Cito, J., Leitner, P., Zdon, U., Gall, H.C.: We're doing it live: an empirical study on continuous experimentation. *J. Inf. Softw. Technol.* (2017, under submission)
114. Spärck Jones, K.: Automatic summarising: the state of the art. *Inf. Process. Manage.* **43**(6), 1449–1481 (2007)
115. Sridhara, G.: Automatic generation of descriptive summary comments for methods in object-oriented programs. PhD thesis, Newark (2012). AAI3499878
116. Sridhara, G., Hill, E., Muppaneni, D., Pollock, L., Vijay-Shanker, K.: Towards automatically generating summary comments for java methods. In: *Proceedings of the International Conference on Automated Software Engineering (ASE)*, pp. 43–52. ACM, Piscataway (2010)
117. Sridhara, G., Pollock, L., Vijay-Shanker, K.: Automatically detecting and describing high level actions within methods. In: *Proceedings of the International Conference on Software Engineering (ICSE)*, pp. 101–110. IEEE, Piscataway (2011)
118. Sridhara, G., Pollock, L., Vijay-Shanker, K.: Generating parameter comments and integrating with method summaries. In: *Proceedings of the International Conference on Program Comprehension (ICPC)*, pp. 71–80. IEEE, Piscataway (2011)
119. Stefan, P., Horky, V., Bulej, L., Tuma, P.: Unit testing performance in java projects: are we there yet? In: *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering, ICPE '17*, pp. 401–412. ACM, New York (2017)
120. Tu, Z., Su, Z., Devanbu, P.: On the localness of software. In: *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2014*, pp. 269–280. ACM, New York (2014)
121. VanHilst, M., Huang, S., Mulcahy, J., Ballantyne, W., Suarez-Rivero, E., Harwood, D.: Measuring effort in a corporate repository. In: *IRI*, pp. 246–252. IEEE Systems, Man, and Cybernetics Society, Piscataway (2011)
122. Vassallo, C., Panichella, S., Di Penta, M., Canfora, G.: Codes: mining source code descriptions from developers discussions. In: *Proceedings of the International Conference on Program Comprehension (ICPC)*, pp. 106–109. ACM, New York (2014)
123. Vassallo, C., Zampetti, F., Romano, D., Beller, M., Panichella, A., Di Penta, M., Zaidman, A.: Continuous delivery practices in a large financial organization. In: *2016 IEEE International Conference on Software Maintenance and Evolution, ICSME 2016*, Raleigh, October 2–7, 2016, pp. 519–528 (2016)
124. Vassallo, C., Schermann, G., Zampetti, F., Romano, D., Leitner, P., Zaidman, A., Di Penta, M., Panichella, S.: A tale of ci build failures: an open source and a financial organization perspective (2017)

125. Vithani, T.: Modeling the mobile application development lifecycle. In: Proceedings of the International MultiConference of Engineers and Computer Scientists 2014, vol. I, IMECS 2014, pp. 596–600 (2014)
126. Wong, E., Yang, J., Tan, L.: Autocomment: mining question and answer sites for automatic comment generation. In: Proceedings of the International Conference on Automated Software Engineering (ASE), pp. 562–567. IEEE, Piscataway (2013)
127. Wu, H.C., Luk, R.W.P., Wong, K.F., Kwok, K.L.: Interpreting tf-idf term weights as making relevance decisions. *ACM Trans. Inf. Syst. (TOIS)* **26**(3), 13 (2008)
128. Xie, T., Pei, J.: Mapo: mining api usages from open source repositories. In: Proceedings of the 2006 International Workshop on Mining Software Repositories, pp. 54–57. ACM, New York (2006)
129. Zagalsky, A., Barzilay, O., Yehudai, A.: Example overflow: using social media for code recommendation. In: Proceedings of the Third International Workshop on Recommendation Systems for Software Engineering, pp. 38–42. IEEE Press, Piscataway (2012)
130. Zhou, Y., Gu, R., Chen, T., Huang, Z., Panichella, S., Gall, H.C.: Analyzing apis documentation and code to detect directive defects. In: Proceedings of the 39th International Conference on Software Engineering, ICSE 2017, Buenos Aires, May 20–28, 2017, pp. 27–37 (2017)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.