Guided pattern mining for API misuse detection by change-based code analysis

Sebastian Nielebock¹ · Robert Heumüller¹ · Kevin Michael Schott¹ · Frank Ortmeier¹

Received: 1 August 2020 / Accepted: 1 August 2021 / Published online: 17 August 2021
© The Author(s) 2021

Abstract

Lack of experience, inadequate documentation, and sub-optimal API design frequently cause developers to make mistakes when re-using third-party implementations. Such *API misuses* can result in unintended behavior, performance losses, or software crashes. Therefore, current research aims to automatically detect such misuses by comparing the way a developer used an API to previously inferred patterns of the correct API usage. While research has made significant progress, these techniques have not yet been adopted in practice. In part, this is due to the lack of a process capable of seamlessly integrating with software development processes. Particularly, existing approaches do not consider how to collect relevant source code samples from which to infer patterns. In fact, an inadequate collection can cause API usage pattern miners to infer irrelevant patterns which leads to false alarms instead of finding true API misuses. In this paper, we target this problem (a) by providing a method that increases the likelihood of finding relevant and true-positive patterns concerning a given set of code changes and agnostic to a concrete static, intra-procedural mining technique and (b) by introducing a concept for just-in-time API misuse detection which analyzes changes at the time of commit. Particularly, we introduce different, lightweight code search and filtering strategies and evaluate them on two real-world API misuse datasets to determine their usefulness in finding relevant intra-procedural API usage patterns. Our main results are (1) commit-based search with subsequent filtering effectively decreases the amount of code to be analyzed, (2) in particular method-level filtering is superior to file-level filtering, (3) project-internal and project-external code search find solutions for different types of misuses and thus are complementary, (4) incorporating prior knowledge of the misused API into the search has a negligible effect.

Keywords API misuses · Error detection · Change-based code analysis · Pattern mining

✉ Sebastian Nielebock
sebastian.nielebock@ovgu.de

Extended author information available on the last page of the article

1 Motivation

Application Programming Interfaces (APIs) enable programmers to reuse existing functionality from libraries. However, since programmers are not always familiar with the particularities of a certain library, they can misuse its API. Generally spoken, these API misuses denote usages that were non-intended by the developers of the library and eventually lead to negative behavior in the client code, for example, performance losses or software crashes.

The reasons why programmers introduce API misuses are manifold and include unknown knowledge on the API usage domain, missing documentation, inconsistent or complex API design (e.g., confusing method names), or unknown internal dependencies of the API implementation (Robillard and Deline 2011; Zibran et al. 2011; Hou and Li 2011; Nadi et al. 2016; Oliveira et al. 2018). In a study on real bug fixes (Zhong and Su 2015), the authors showed that half of their analyzed fixes involved at least one API-specific change to resolve the bug. Even worse, API misuses appear in different shapes. A study of 90 API misuses, therefore, introduced a first Misuse Classification scheme (MuC Amann et al. (2018)). The MuC distinguishes between missing and redundant API elements such as method calls, conditions, iterations, and exception handling. In addition to this classification, other typical classes of misuses are incorrect ordering of method calls or incorrect usage of parameters (Robillard et al. 2013; Frolin S. Ocariza et al. 2013).

Therefore, recent research strives to invent automated methods for detecting API misuses. Due to increasing computational power, data mining techniques have become prominent in finding so-called API specifications (Robillard et al. 2013). An API specification is a formal model describing properties of the correct usage of APIs and thus is used as an oracle to detect API misuses as code parts violating this specification. On the highest level, API specifications have so far been distinguished into two categories, namely, dynamic invariants together with pre- and post-conditions as well as usage patterns (Ammons et al. 2002). Dynamic invariants reason about program state and how it changes with regard to API usages (Ernst et al. 2001, 2007). They vary in their representation form regular expressions over finite-state automata (Ammons et al. 2002; Yang and Evans 2004; Gabel and Su 2008; Pradel and Gross 2009), up to temporal specifications (Wasylkowski et al. 2007; Wasylkowski and Zeller 2011). Similarly, various representations of API usage patterns have been proposed in prior research such as API method call pairs (Weimer and Necula 2005), association rules (Livshits and Zimmermann 2005; Li and Zhou 2005), API method call sequences (Thummalapenta and Xie 2007), trees (Allamanis and Sutton 2014), or graphs (Nguyen et al. 2009; Amann 2018). Some approaches use Bayesian inference to learn the correct usage of APIs from code examples as probability distribution (Allamanis and Sutton 2014; Murali et al. 2017). Zhou et al. infer specifications from the API documentation to detect inconsistencies with the respective specifications from the code (Zhou et al. 2017). In essence, they represent structural or temporal constraints between the elements of an API.

In this paper, we focus on API usage patterns that are inferred from existing source code through data mining. Moreover, we only consider intra-procedural patterns, namely, those that only occur within single method declarations. Thus, our results do not refer to inter-procedural patterns which are scattered among multiple method declarations. The general procedure of mining usage patterns and detecting API misuses comprises the following five steps:

1. Collect a representative set of source code for mining
2. Transform this code set into an intermediate representation (e.g., execution traces (Yang et al. 2006), syntax trees (Allamanis and Sutton 2014), API usage graphs (Amann 2018))
3. Conduct a frequent pattern mining approach (e.g., association rule mining (Li and Zhou 2005), sequence mining (Zhong et al. 2009), subgraph mining (Amann 2018)) on this representation
4. Filter generated patterns based on suitable ranking metrics (e.g., support, confidence, or others (Le and Lo 2015))
5. Compare the usage of the API with those of the highest-ranked patterns and report violations as misuses (Amann 2018)

Research on API misuse detection and API usage pattern mining, in particular, has focused on reducing the high number of false positives, i.e., patterns originating from random co-occurrences of code elements, (e.g., method calls). These false positives can cause false alarms during API misuse detection which impedes the practical application of such automated detectors. Particularly, the approaches mainly improve the last four steps, namely, the intermediate representation, the frequent pattern mining approach, the filtering and ranking strategy, as well as the violation detection (Li and Zhou 2005; Livshits and Zimmermann 2005; Thummalapenta and Xie 2007; Wasylkowski et al. 2007; Gabel and Su 2008; Nguyen et al. 2009; Pradel and Gross 2009; Zhong et al. 2009; Wasylkowski and Zeller 2011; Allamanis and Sutton 2014; Murali et al. 2017; Amann 2018).

However, little effort (e.g., by Le Goues and Weimer (Le Goues and Weimer 2012)) was put into the initial step of collecting code before mining, even though the quality of the input data has a significant impact on the results of any data mining algorithm. False or noisy data results in bad—or at least unpredictable—results, as seen with classifiers (Agrawal and Menzies 2018). Moreover, existing research gives only little insight into how code collection as well as subsequent pattern mining and misuse detection could conceptually be included in a software development cycle.

For that purpose, this paper concentrates on the investigation of strategies for collecting code samples before the mining step. Our envisioned strategies aim to improve the true positive rate agnostic of the concrete static, intra-procedural mining tool. In particular, such a strategy should select source code that contains a high density of relevant patterns concerning a particular API misuse. We refer to such patterns as *fixing patterns*, meaning patterns that can detect and fix an API misuse. Thus, by increasing the *relative frequency* of fixing patterns in the

data set, we increase the likelihood that support-based miners will find the fixing patterns. Note that we do not directly target the goal of decreasing the false positive rate since this was mainly part of previous API misuse detectors. However, our investigated search strategy gives an additional lightweight pre-processing step that improves the performance (i.e., the true positive rate) by reducing the number of code samples to mine from.

Moreover, we introduce a concept for how this strategy can be incorporated into a standard software development cycle. We exemplarily show that, due to the lightweight design of our search process, it requires only a little additional effort compared to later mining and filtering steps.

Our collection strategy is based on the analysis of code changes, i.e., commits in a version control system. The idea is to incrementally analyze only the small subsets of code that are affected by a change and to use this very specific context for a focused API misuse detection. To this end, we search for source code files with similar but correct API usages regarding the changed code and further filter these using different strategies. We compared these strategies by using known, real-world API misuses and their respected fixes from the two benchmarks MUBench (Amann et al. 2016) and the AU500 (Kang and Lo 2021). In particular, we determine how frequently the known (or similar) fixes are found in the filtered sets and use this information to identify the best strategy (i.e., that with the highest relative frequency).

Afterward, we check whether the results of an existing support-based API pattern mining approach are actually improved by comparing the performance with and without the filtering strategy.

This way, we answer the following three research questions (RQs):

- RQ₁ Does the changed-based code analysis sufficiently reduce the number of code snippets to efficiently perform API usage pattern mining
- RQ₂ Which filtering strategy yields to the highest relative frequency of fixing patterns in the retrieved source files?
- RQ₃ Does an existing API usage pattern miner increase the number of detected fixing patterns by means of the selected filtering strategy?

For replicability, we provide our data sets and evaluation scripts as a replication package¹. This package can also be seen as a first prototype of the search and mining process. For example, this could be installed on a standard continuous integration (CI) server similarly as, e.g., automated testing approaches. Currently, a developer who commits her code changes to such a CI system would receive API usage patterns similar to the API usages in their change. Potential subsequent steps, such as misuse detectors and misuse fixing approaches, could also be added to the CI system. This way, misuse detection and correction are directly applicable for software developers.

¹ available at <http://doi.org/10.5281/zenodo.5091822>.

Fig. 1 Envisioned API misuse detection and correction process

Our process and the analyzed filtering strategies are introduced in the following Sect. 2. Then, we present our results of the evaluation of the three research questions (Sect. 3). Afterward, we discuss potential threats to validity (Sect. 4) of our results as well as differences and similarities to related work (Sect. 5). Finally, we conclude our results and present future work (Sect. 6).

2 Process and strategies of a change-based API misuse detection

Within this section, we first present our envisioned process of an API misuse detection which leverages code search and filter strategies to improve the input data for API usage pattern mining. This section describes the concrete use case for which the collection strategies and subsequent mining approaches are designed. Thus, the misuse detection is not part of the contribution of this paper. Second, we discuss the notion of different search and filter strategies.

2.1 A vision of a change-based API misuse detection and correction

Figure 1 depicts our envisioned API misuse detection process. It describes how our analyzed search and filter strategies, the API usage pattern mining, as well as the final detection and correction, can be seamlessly integrated into an ordinary continuous integration (CI) process.

We consider a developer who commits changes in her client project into a version control system. Assuming that some of these changes could contain an API misuse (①), our process conducts an *API change analysis* (②) based on the changeset of that commit. The goal of this step is to locate the changed methods and to extract the affected API usages and their context from each method. Based on this information, step ③ searches for each changed method other source code samples (either from

this or from foreign repositories) with similar API usages. These, in turn, are further filtered (④). The steps ③ and ④ describe the search and filter strategies, which we present in detail in the upcoming section and evaluate in Sect. 3.

The filtered code is then transformed (⑤) into an intermediate data representation (e.g., method call sequences, abstract syntax trees, API usage graphs) to be further processed by an API usage pattern miner. This miner conducts a frequent pattern mining approach (⑥) and generates a list of ranked API usage patterns (e.g., ordered by support).

Then, the process checks whether the changed version of the client code violates one or multiple of the highest-ranked API usage patterns (⑦). In case of violations, one can generate patch candidates based on the violated usage patterns (⑧), for example, by using the difference between pattern and code as a patch. After validating the misuse and selecting a fitting patch, the developer can apply the fix (⑨).

We envision these steps to be set up in an ordinary CI process. Thus, whenever developers commit code changes they get instant feedback whether or not the changes introduced an API misuse. In case of a misuse, commits can be automatically requested to be revised, for instance, by the patches suggested to correct the misuse.

Note that this approach assumes that every commit contains a potential API misuse. Therefore, the API change analysis needs to reduce the number of methods to be analyzed so that the number of subsequent mining runs remains low as well. Moreover, to reduce the effort for a single mining run, the search and filter strategies effectively reduce the number of similar source files without harming the true-positive rate of the miner. However, in case the subsequent pattern mining requires huge computation time, one may limit this analysis to specific testing branches in the CI system, or require a developer to deliberately trigger it.

This process shares some similarities with the work by Saied et al. (2020), which integrates API misuse detection as an interactive element in the coding task. However, our process does not require storing a set of previously inferred patterns but conducts online pattern inference during analysis from an evolving code base. Since here, the preprocessing and mining require some amount of time (approx. 5–10 min) it is currently neither intended nor recommended for interactive usage.

2.2 Search and filtering strategies

To improve the input data of the mining step, we propose different strategies for searching and filtering source code to increase the ratio of relevant source code in terms of finding fixing patterns. In this section, we describe the main steps and the intuition behind these strategies. Note that our evaluation (cf. Sect. 3) is based on Java as a programming language. Therefore, some technical details refer to particularities of that language and may need to be adapted for use with other programming languages.

In the general process (cf. Fig. 2), first, we extract the set of changed methods (Ⓐ) from a single commit. Then, for each changed method, we conduct a separate API change analysis (Ⓑ). This analysis extracts a set of relevant API import

Fig. 2 Process steps of the search and filtering strategies

statements, i.e., those that import a type of a third-party library used in the analyzed method. It further extracts a set of *keywords* describing the context of the API usage, i.e., within the analyzed method. Particularly, these keywords are the set of included class names from third-party libraries, as well as all method names used within the analyzed method. Using the set of API import statements, we conducted a code search for files that also import these types (C). Afterward, we filtered the files (D) and the method declarations (E). In the following paragraphs, we describe each step in detail.

Commit The commit step (A) extracts the set of changed methods from a commit that may have introduced API misuses. Note that during our evaluation, we specifically investigate misuse-introducing commits of known API misuses. We obtained these commits from the information given in the analyzed benchmark. Details are described in Sect. 3.

We inferred the changed methods using the version control system (i.e., `git diff`) and extracted the set of changed source files with the respective changed lines. Then we parsed these changes and located those method declarations that were at least partially affected by these changes. These methods constitute the scope for the following analyses. We cannot restrict these only to the changed lines, because it does not necessarily contain all information required to detect misuses. For example, based on a previous method call order, an additional method call may introduce a misuse, e.g., an invalid double initialization of an object. Moreover, the effort to analyze the method scope is still manageable. This is important as a complex and long-lasting analysis would impede the development process.

API change analysis For each method detected in the previous step, we initiate a separate API change analysis and a subsequent search and filtering process. In the API change analysis (B), we want to discover which API elements from outside of the current project's scope were changed by the commit.

Note that we only considered third-party libraries for two reasons. First, for project-internal API elements, i.e., types and methods that are declared within the analyzed project, it is very unlikely to find usage patterns in external code. Since we are comparing project-internal and project-external (i.e., in foreign projects) code search strategies this comparison would be fairly biased. Second, usages of the `java.lang` APIs are far too common and introduce too much noise into the filter process.

Fig. 3 Example of a keyword extraction for the `doSomething`-method. Left: Source Code from which keywords are extracted. Right: Set of extracted API import statements and keywords

Intuitively, the discovered API elements correspond to the keywords that a developer would use when searching for similar code on the web. To identify useful API elements, we reviewed the real-world misuses from the MUBench benchmark (Amann et al. 2016). Based on these code samples, we identified common patterns of code features that describe the API usage and its context. In addition, Zhong et al. provided some insights into which code features indicate API usage (Zhong et al. 2009).

We exemplify the identified code features describing the relevant API elements in Fig. 3.

The used data types of third-party libraries are important indicators for API usage. In general, we consider a data type as a relevant API element if it is used in the analyzed method and originates from a third-party-library, i.e. is explicitly imported via an `import` statement and does not originate from the project itself. The usage of a data type in a method means one of the following five alternatives:

- The type is applied as parameter type of that method
- The type is applied as a return type of that method
- The type is applied as `thrown` type in the method's declaration
- The type is explicitly mentioned in an expression inside of the method's body
- The type is inherited from the class of that method and the method overrides the method declaration, explicitly denoted by an `@Override` annotation

Considering the method `doSomething` in our example, this covers the data types `AClass`, `BClass`, and `ZClass`. Since type `CClass` is not used in the method, it is not extracted.

We consider data-types as imported when there exists an explicit `import`-statement that ends with that type name. For example in Fig. 3, the type `BClass` used in line 12 is imported via the `import`-statement in line 4 (`import a.b.BClass;`). On the other hand, assume that the type `RClass` used in line 14 is imported via the wildcard `import`-statement in line 7 (`import x.v.*;`). Since this type name is not explicitly mentioned, it is not considered as relevant.

Moreover, we checked whether the used types were related to a project-internal or a third-party library. This was done by checking whether the `import`-statements have the same prefix (i.e., qualifiers of that data type) as the `package`-statement in that class declaration. Particularly, we checked whether the first three qualifiers are identical to those of a particular import (e.g., `my.own.pkg` in Fig. 3). The rationale is that regarding the naming convention of packages in Java², the first three qualifiers usually identify the package. In case that only one or two qualifiers are used in the package, we only check whether these are prefixes of the respective `import`-statement. If no package name is given, which did not occur in our evaluation data set, we ignore all data types. For example, the `import`-statement of type `QClass` in line 8 has the same prefix, i.e. `my.own.pkg`, as the respective `package`-statement. Thus, it is not extracted as a relevant API element.

All `import`-statements whose data types were found to be relevant API elements are added to the *API imports set*. The respective type names are added to the *keywords set*.

Note that besides the data types used in the method declaration and its body, we also add the names of inherited types, if the method under investigation is overridden (indicated by the `@Override` annotation). In such a case, we also add the method name of the overridden method declaration, which is usually not considered, to the keywords set. This is reasonable since some framework APIs (e.g., Eclipse, Android) are frequently accessed by inheritance and thus could be misused.

Additionally, we extended the keyword list with the names of all methods that were called in the investigated method. This included calls to project-internal and `java.lang` methods. The reason is that we only consider source code changes and we cannot completely resolve all data types from this partial code. Thus, sometimes it was not possible to decide whether a method belonged to an internal, `java.lang`, or a third-party-library type.

Searching source files ($search_{loc}$ and $search_{imp}$) The goal of the code search is to find code that is similar to the investigated method in terms of its API usages. In our cascaded search process, the first step (C) applies the *API import-set* as a set of keywords to find files that imported the relevant API types that were identified in the previous step.

Here, we used the two different search strategies $search_{loc}$ and $search_{imp}$. First, we distinguished between where we searched for similar code, namely *internal* code, i.e. from the same project, or *external* code, i.e. originating from a foreign project ($search_{loc}$). Second, we varied which *API imports* were used. In the first version, we applied *all extracted import statements*. In the second one, we only used the *misused*

² <https://docs.oracle.com/javase/tutorial/java/package/namingpkgs.html>.

import statements, i.e., statements that imported misused APIs ($search_{imp}$). We discuss the rationale for both strategies in the following.

Regarding $search_{loc}$, Amann discussed that other API usages can be found either *internally*, i.e. in the same project, or *externally*, i.e. in other projects (Amann 2018). The internal search can find correct API usages or already fixed API misuses in other locations of the same project. This is, for example, indicated by the plastic surgery hypothesis from automatic program repair (Le Goues et al. 2019). On the other hand, the external code search is likely to provide more and diverse data, and thus increases the likelihood to find similar source code.

We applied $search_{imp}$ using different sets of extracted import statements, namely, *all* vs. *misused* imports. We assume that searching with misused imports will increase the true positive rate. However, it also requires a preliminary analysis to extract them since, in practice, we usually don't know the misused API. Therefore, we check whether such a preliminary analysis would significantly increase the likelihood to find fixing patterns or not. Note that in our evaluations, we already know the misused API from the information in our data set (cf. Sect. 3.1) and therefore we did not implement such an approach.

File filtering ($filter_{file}$) After searching, our process filters files regarding the keyword set (D) that is further denoted as $filter_{file}$. We do not expect that all keywords have to be used in a similar file because some keywords could belong to a co-applied API or method names of internal APIs. Therefore, we introduce a measurement—so-called *satisfaction ratio* (*sr*) for files to estimate to which degree these contain the keywords. The satisfaction ratio describes the proportion of keywords found in a source file $srcFile$ with respect to the keyword set $kwSet$. It is defined as follows:

$$sr(srcFile, kwSet) = \frac{|\{kw \in kwSet \text{ if } srcFile \text{ contains } kw\}|}{|kwSet|}$$

A satisfaction ratio of 0 does not require *any* keywords to be present, while a satisfaction ratio of 1 requires *all* keywords to be present. We do not believe that either extreme is useful to increase the relative frequency of fixing patterns. While the first strategy does not require the file to be similar at all, the second strategy would yield too few results. Thus, in our evaluation, we consider a range of values between 0 and 1 for the satisfaction ratio.

Note that this step could be technically integrated into the previous code search. However, in this paper, we want to determine the effect of each single filtering step and therefore these steps are separated.

Method filtering ($filter_{method}$) In the second, so-called $filter_{method}$ strategy (E), we extracted those methods from the source files that contained at least one of the keywords in the keyword set. This is a more fine-grained approach. Moreover, the envisioned API misuse detection considers the method scope, and therefore reducing the number of methods while keeping related ones can increase the relative frequency of the patterns.

We filtered the methods by parsing each file and generating the token sets for all methods. After removing syntax elements and Java keywords, we check if at least one keyword is contained in the remaining set.

Note that we did not apply several different satisfaction ratios as in the previous step. This decision was based on three reasons. First, this would further increase the number of strategy configurations to be analyzed. As shown in Sect. 3.3, the current number of configurations to be tested per misuse is 40 ($2 \cdot 2 \cdot 5 \cdot 2$) eventually leading with our 37 misuses from the MUBench dataset to 1480 different configurations. By adding an equal fragmentation with the five different *sr*-values and assuming that we do not have to run method filtering with higher *sr* on method level than on file level (e.g., a file filtered out with a *sr* = 0.5 certainly does not contain a method with *sr* = 0.75) the number of configuration would increase to 60 per misuse ($2 \cdot 2 \cdot (\sum_{i=1}^5 i)$). Thus we would have to analyze at most 2220 configurations. Second, we can assume that the optimal *sr* on the method level is usually lower than the one for the file level since method declarations contain fewer keywords. Therefore, the range to be analyzed must be lower, for instance, the interval [0, 0.25]. However, then we need to test even more configurations since then we cannot exclude certain configuration combinations. Third, we deemed the effect to be non-significant compared to the effort of analyzing further configurations (i.e., generating AUGs and testing the pattern containment). Nevertheless, we computed the average *sr* for the filtered method (i.e., those that contain at least one keyword). These values can then be used to further improve the method filtering.

After introducing the filtering steps, in the upcoming section, we discuss how we evaluate our process and in particular the different strategies for searching and filtering source files. We then interpret the results of the evaluation for our three research questions.

3 Evaluation

We primarily evaluated our approach by means of the MUBench benchmark (Amann et al. 2016)³. This benchmark represents a set of real, validated API misuses from open-source projects. Since we found not all misuses to be suitable for our evaluation, we selected a subset for our analyses. We describe this data acquisition in Sect. 3.1. For RQ3, we additionally incorporated the AU500 dataset⁴ by Kang et al. consisting of 500 API usages manually labeled correct and incorrect API usages (Kang and Lo 2021). Moreover, we explain how we obtained the API misuse-introducing commits and the respective *API Usage Graphs*, the intermediate data representation, based on the implementation by Amann (2018)⁵. Afterward, we evaluate our three research questions. For each question, we first, describe our methodology, second, present the results of our evaluation, and third, summarize the main results and implications.

³ <https://github.com/stg-tud/MUBench>.

⁴ AU500 is part of the ALP software artifact at <https://github.com/ALP-active-miner/ALP>.

⁵ <https://github.com/stg-tud/MUDetect>.

3.1 Data acquisition and processing

Selecting API misuses We considered an initial set of 245 API misuses from Java projects obtained on February, 5th 2019 from the MUBench benchmark⁶ (Amann et al. 2016). For each misuse in this benchmark, the authors provide a file describing the meta information of the misuse, which essentially includes the version control system, the misused API, and the fixing commit (i.e., that commit that fixed the misuse). From these misuses, we selected those 103 that originated from projects using the git-version control system and for which the benchmark provides a fixing commit. The rationale for git is that it is one of the most frequently applied version control systems⁷. Since we need the fixing commit to identify the misuse-introducing commit, this is also a mandatory requirement.

Then, we further removed 66 misuses due to the following reasons:

- Misuses are essentially duplicates, i.e., the misuse of the serialization of an object in a testing context in the jodatime-project—we only kept one version of that misuse (36)
- Misuses of a `java.lang-API`, which is not covered by our method (18)
- Misuses of an internal, i.e., project-related, API, from which we do not expect to find correct usages in external projects (8)
- Non-distinguishable misuse, i.e., same API misuse in the same method in the same class—we only kept one version of each misuse (2)
- Misuse is a false parameter, which cannot be represented by the used intermediate representation (2)

Thus, we kept 37 misuses for our analysis.

The AU500 consists of 500 manually labeled API usages from 16 open-source projects (Kang and Lo 2021). These projects are disjunct from those of the MUBench dataset, and thus, this dataset is well suited for an independent validation. 385 entries of this dataset are labeled as correct while the other 115 are marked as misuses. All entries have the meta-information on the git-repository, the commit hash of that version, and the source file as well as the method containing the API usage. We use all those usages for the analysis of RQ3.

Detecting API misuse-introducing commits In MUBench its creators already identified the fixing commit, i.e. the commit in which the misuse was corrected. However, we are also interested in the *API misuse-introducing commit*, i.e., the commit that made the changes that lead to the misuse. For that purpose, we checked out the fixed version and identified those lines of code that had to be changed to fix the misuse via the command `git diff`. We then obtained the previous version of the fixing commit and run the command `git blame`⁸ to identify in which commit these lines were added to the repository. We denote this commit as the misuse-introducing commit. Note that in the case of multiple different commits, i.e., among

⁶ commit b8124077 from <https://github.com/stg-tud/MUBench.git>.

⁷ <https://www.openhub.net/repositories/compare>.

⁸ <https://git-scm.com/docs/git-blame>.

differently added lines, we chose the latest commit, since this indicates the point in time when the misuse was ‘complete’. This essentially is a git-adapted version of the SZZ-algorithm (Śliwerski et al. 2005), which was designed for usage with the CVS version control system. With these remaining 37 misuse-introducing commits, we evaluated the change-based code analysis.

In the AU500 dataset, not all usages are misuses. This is why we do not determine the misuse-introducing commits for this dataset. In contrast, we extracted the keywords and import statements only from the single revision of that API usage. We identified this revision by the respective commit hash, source file, and method name. Particularly, we checked out the respective commit and analyzed the complete method declaration as if it was added all at once using the API Change Analysis as described in Sect. 2.2.

Collecting similar source files We also need different sets of source files for the *internal* and *external* code search (cf. Sect. 2.2). For the *internal search* we used all source files from the same project and revision of the misuse-introducing commit, excluding the file that contains the misuse.

We conducted the *external search* by means of the *Searchcode* engine⁹. Searchcode accesses well-known code repositories such as *GitHub*, *BitBucket*, *Google Code*, and *GitLab*. Compared to other code search engines such as *Boa* (Dyer et al. 2013) and *GHTorrent* (Gousios 2013), this engine provides access to individual source files without having to download the whole project. Due to internal restrictions at the time of our analysis, Searchcode returns at most 1,000 source files, which are ordered by relevance. We contacted Searchcode’s developer to clarify the definition of relevance and were informed that it is estimated based on the proximity of detected keywords in a file. Thus, a file containing the keywords as “foo bar” is ranked higher than a file having “foo” and “bar” separated in different parts of the file. We ran two search sessions for each misuse to collect source files. In the first one, we only searched with the misused API import statement(s). In the second session, we used all extracted import statements. We downloaded both sets via Searchcode’s REST API¹⁰ between February 7th and February 8th 2019 and eventually merged both sets, yielding up to 2,000 source files. Due to an error, we repeated the analysis for the `logblock-logblock-2_15` misuse on December 12th 2019. For the analysis of the AU500 dataset, we downloaded the source files on June 15th 2021. We prevented source files from the same project from being found by excluding all files with the same prefix in the `package`-statement as used in that original project. Moreover, we excluded all source files whose generation of the API Usage Graphs occupied too much memory and caused the generation script to crash on our evaluation system. In particular, we had to exclude externally found source files for 13 misuses (i.e., nine misuses with a single excluded file and four misuses with two up to nine excluded files) for the MUBench dataset.

Note that we kept both code file sets (i.e., internal and external without exclusions) consistent for all subsequent search and filtering steps. Therefore, the

⁹ <https://searchcode.com/>.

¹⁰ <https://searchcode.com/api/>.

Fig. 4 Example of an API usage graph of the `myFancyMethod`-method in the `SampleClass` on the left hand-side

potential bias introduced by the search algorithm of Searchcode is also consistent across all search and filter strategies.

API usage graphs For the analysis of RQ_2 and RQ_3 , we utilized an intermediate source code representation, namely the API usage graph (AUG) introduced by Amann (2018). This directed, labeled multigraph is a static code representation, which depicts data- as well as control flow properties. In this respect, it is specifically tailored to representing the API usages of a single method and is therefore ideal for our analyses. Moreover, this data structure enables us to also reuse the corresponding API usage pattern miner, which was also introduced by Amann.

Figure 4 shows a small example of an AUG. The AUG consists of different types of nodes and edges. Rectangles denote *action* nodes, e.g. method calls, while ellipses represent *data* nodes, e.g. object instances. In addition to simple method calls, e.g. `doSomething`, there also exist special actions, e.g. `<init>` for constructor calls or `<return>` for return-statements. Data nodes usually represent object instances labeled with their respective type. If a type cannot be inferred statically from the code it is labeled as ‘UNKNOWN’. In the example, the return type of the `doSomething` method is UNKNOWN since the declaration of the method is missing and therefore the type resolution could not decide whether it is of type `Integer` or one of its subclasses. Besides having different types of nodes, AUGs also feature different kinds of edges, i.e., control flow edges (dashed) and data flow edges (solid). Control flow edges describe structural properties of the code, e.g. the order of actions (*order*-edge), while data flow edges describe how information in the form of objects is processed through the code. This includes instance definition (*def*-edge), calling methods on instances (*recv*-edge), or using instances as parameters to other methods (*para*-edge).

For further details on the AUG and the miner please refer to the work of Sven Amann (2018).

3.2 Commit size analysis (RQ₁)

Methodology First, our approach uses commits to reduce the amount of source code that is analyzed regarding an API misuse. In this first experiment, we analyze the typical size of commits that contain API misuses, i.e., their number of changed methods, relevant external API imports, and their number of extracted keywords. Since API usage pattern mining may require a lot of computational power and memory consumption, these values can be crucial for the overall performance. For example, if a commit changes many methods, we also have to conduct the same number of searching and mining tasks. Therefore, a small number of analyzed methods per commit is desirable.

We investigated these characteristics using the 37 misuse-introducing commits obtained from MUBench in the previous step. For each commit, we conducted the API change analysis by first determining the number of changed methods, and second, extracting for each changed method the set of API import statements and keywords as presented in Sect. 2.2. For parsing source code, we used Eclipse's JDT parser¹¹.

We locally stored each changed method. To avoid name clashes for methods (e.g., caused by overloading), we added a numerical ID to each method. For every method, we then stored the sets of API imports and keywords. We analyzed all collected information via python scripts and provide both the data and the evaluation scripts in our replication package¹.

Results In Table 1, we show detailed information on each misuse and its respective misuse-introducing commit. In addition, it contains the number of all methods in the project (Column A), the number of methods changed in the misuse introducing commit (Column C), and the subset of those methods that were part of an external API (Column E). Note that for Column A, we obtained the total number of methods by parsing only unique source files identified by their md5-hash value. Thus, two methods originating from two identical source files are only counted once. Moreover, some misuses were introduced in the same commit. Therefore, we analyze the degree of method reduction only for the 32 *unique commits*.

First, we considered all methods that were affected by the commit, regardless of whether this change edited an external API or not. Figure 5 plots the distribution of misuse-introducing commits for increasing numbers of changed methods. The majority of 25 misuses modified less than 100 methods, while nine outliers had up to 2517 changed methods (i.e., misuse `bcel_101`). When considering only those methods for which we found imports of third-party libraries, the huge numbers shrink drastically as denoted by Fig. 6. Since we only consider API misuses of third-party libraries, we do not need to investigate methods for which we cannot infer an

¹¹ <https://www.eclipse.org/jdt/>.

Table 1 Misuses with the characteristics of their misuse-introducing commits

#	Misuse	Repository (subdomain at https://github.com)	MIC	#Methods			Reduction (%)	
				A	C	E	A2C	C2E
1	alibaba-druid_1	/alibaba/druid.git	de13143e0	16095	12	8	99.9	33.3
2	alibaba-druid_2	/alibaba/druid.git	de13143e0	16095	12	8	99.9	33.3
3	android-rs-resjta_1	/android-rs-resjta.git	b3445d9	10817	2275	642	79.0	71.8
4	androiduil_1	/nostra13/Android-Universal-Image-Loader.git	9d77de9	737	279	131	62.1	53.0
5	apache-gora_56_1	/apache/gora.git	e4db20a	1565	141	57	91.0	59.6
6	apache-gora_56_2	/apache/gora.git	bbad5d213	1424	39	35	97.3	10.3
7	bcel_101	/apache/commons-bcel.git	d532ec1	3475	2517	269	27.6	89.3
8	calligraphy_1	/chrisjennx/Calligraphy.git	1a2d0f5d	32	10	8	68.8	20.0
9	calligraphy_2	/chrisjennx/Calligraphy.git	1a2d0f5d	32	10	8	68.8	20.0
10	closure_2	/google/closure-compiler.git	e5d3e5e012	11135	28	14	99.7	50.0
11	jodatime_269	/emopers/joda-time.git	08a925a31	4429	54	10	98.8	81.5
12	jodatime_339	/emopers/joda-time.git	9b01b9e8b	9054	21	11	99.8	47.6
13	jodatime_361	/emopers/joda-time.git	7fe68f297	2556	2451	519	4.1	78.8
14	jodatime_362	/emopers/joda-time.git	7fe68f297	2556	2451	519	4.1	78.8
15	jodatime_363	/emopers/joda-time.git	7fe68f297	2556	2451	519	4.1	78.8
16	lnreadera_1	/calvinaquino/LNReader-Android.git	a514f35	3329	81	72	97.6	11.1
17	lnreadera_2	/calvinaquino/LNReader-Android.git	a514f35d	3329	81	72	97.6	11.1
18	logblock-logblock-2_15	/emopers/LogBlock-2.git	5ealb0b	70	5	4	92.9	20.0
19	mqtt_389	/emopers/paho.mqtt.java.git	77aa39b9	670	608	115	9.3	81.1
20	mqtt_390	/emopers/paho.mqtt.java.git	f60b3721	990	59	23	94.0	61.0
21	onosendai_1	/haku/Onosendai.git	Cf2de97	1618	3	2	99.8	33.3
22	openiab_1	/onepf/OpenIAB.git	00e5612	957	173	100	81.9	42.2
23	screen-notifications_1	/lkorth/screen-notifications.git	fa75a61f	48	21	19	56.2	9.5
24	tbktu-ntru_473	/emopers/ntru.git	4a095cc	399	13	5	96.7	61.5

Table 1 (continued)

#	Misuse	Repository (subdomain at https://github.com)	MIC	#Methods			Reduction (%)	
				A	C	E	A2C	C2E
25	tbuktu-ntru_474	/emopers/ntru.git	e4f8688	187	8	4	95.7	50.0
26	tbuktu-ntru_475	/emopers/ntru.git	8cb6471	521	41	17	92.1	58.5
27	testing_16	/cbeust/testing.git	234c85874	5557	21	19	99.6	9.5
28	testing_17	/cbeust/testing.git	b68cf6de8	5479	18	17	99.7	5.6
29	testing_21	/cbeust/testing.git	24341340b	5432	14	11	99.7	21.4
30	testing_22	/cbeust/testing.git	79cd443f	4395	4	2	99.9	50.0
31	thebluealliancea_1	/Adam8234/the-blue-alliance-android.git	be7b752	1168	10	10	99.1	0.0
32	thomas-s-b-visualee_29	/emopers/visualee.git	14c3f03	152	76	33	50.0	56.6
33	thomas-s-b-visualee_30	/emopers/visualee.git	14c3f03b	152	76	33	50.0	56.6
34	thomas-s-b-visualee_32	/emopers/visualee.git	d4dc0ba	250	1	1	99.6	0.0
35	tucanmobile_1	/Tyde/TuCanMobile.git	805f770	62	11	9	82.3	18.2
36	ushahidia_1	/ushahidi/Ushahidi_Android.git	db2b310	4405	63	40	98.6	36.5
37	wordpressa_1	/wordpress-mobile/WordPress-Android.git	88368deadbe	5453	70	39	98.7	44.3

MIC misuse-introducing commit, *A* all methods, *C* changed methods in the MIC

E All methods from C that contain at least one external (third-party) API

A2C Reduction from all to changed methods

C2E Reduction from changed to changed methods that contain at least one external (third-party) API

Fig. 5 Distribution of misuse-introducing commits among the number of changed methods (bin size of ten)

Fig. 6 Distribution of misuse-introducing commits among the number of changed methods with at least one extracted external API (bin size of ten)

import statement. Then, 26 misuses changed less than 100 methods, with 18 of them having less than 20 changed methods. Still, there exist six extreme outlier commits with 100 or more changed methods.

Our results show that we can effectively reduce the number of methods for later API misuse detection by considering only changed methods. Particularly, we compared the reduction against the number of methods that existed in the commit right after the misuse-introducing commit as a starting point. Considering all 32 unique commits we have an average reduction of 81.9% (median 96.2%). Further, we decreased the number of methods by checking whether these contained an API from a third-party library. By that means, we further reduced the number of methods on average by 15.6% (median 11.6%). An interesting observation is that in cases in which the change-based approach could not reduce many methods (i.e.,

Fig. 7 Distribution of the number import statements among the misuses for methods with at least one third-party import involved

Fig. 8 Distribution of the number extracted keywords among the misuses for methods with at least one third-party import involved

jodaetime_361-jodaetime_363 and mqtt_389) this step could effectively do so. The mean number of changed methods is ≈ 287.6 (median 33.5), while after

Table 2 Different configurations for the analysis of file searching (i.e., $search_{loc}$ and $search_{imp}$) and filtering (i.e., $filter_{file}$ and method $filter_{method}$)

$search_{loc}$	Internal			External		
$search_{imp}$	All imports			Misused imports		
$filter_{file}$	sr = 0.0	sr = 0.25	sr = 0.5		sr = 0.75	sr = 1.0
$filter_{method}$	Applied			Not applied		

the removal of methods without a change to a third-party library this number is reduced to ≈ 71.2 (median 18). In total, we could reduce the number of methods on average by 86.4% (median 96.8%).

In a second step, we investigated the number of extracted import statements and the number of keywords for those methods that referenced at least one external API (i.e., have at least one extracted import statement). These values are interesting since they indicate how many APIs potentially have to be analyzed. Moreover, having a huge number of keywords would also reduce the number of files satisfying the *satisfaction ratio* in the *file filtering* step. On the other hand, it increases the chance of including more methods in the *method filtering* step, since more methods may match at least one of these words.

Figure 7 shows the distributions of the numbers of import statements among the changed methods with at least one external API for each misuse. We can observe that the majority of methods have at most 5 import statements estimated by the upper whiskers (1.5 of the interquartile range). On average, for methods with at least one imported third-party API, 1.8 (median 1) import statements were found. However, there are still outliers that refer up to 28 imports (e.g., `android-rcs-rcsjita_1`).

Regarding the number of keywords (Fig. 8), usually, at most 20 keywords are extracted - once again estimated by the upper whiskers. In extreme cases (e.g., `android-rcs-rcsjita_1`), the number of keywords rises up to 79. On average, 6.3 keywords are extracted (median 4).

Moreover, we checked the ability of the code change analysis to extract the import statement(s) of the misused API(s) in the investigated method. This is important since our method uses these imports to find similar API usages. In case we did not extract the misused API, we hardly expect to find similar API usages for that misuse. We know the misused API(s) based on the meta-description in the MUBench dataset. For 31 out of 37 misuses, our method successfully extracted the import statement of the misused API.

We did not precisely measure the execution time for a single change-based analysis. However, we can recall from the timestamps of the generated files, that a single execution takes from several seconds up to at most three minutes for large commits (e.g., `android-rcs-rcsjita_1`). Note that this time excludes the downloading of the repository source files.

Implications Our results show that the change-based API analysis can effectively reduce the amount of source code that has to be analyzed. At the same time, it is still able to determine the respective misused API in 31 out of 37 cases and on average

Table 3 Overview under which conditions the groups for statistical comparison of the single strategies were obtained

Comparison of	Condition to obtain independent groups
$search_{loc}$	no $filter_{file}$ (i.e., $sr = 0$); no $filter_{method}$; each single $search_{imp}$ -strategy
$search_{imp}$	no $filter_{file}$ (i.e., $sr = 0$); no $filter_{method}$; each single $search_{loc}$ -strategy
$filter_{file}$	no $filter_{method}$; each single $search_{imp}$ and $search_{loc}$ strategy
$filter_{method}$	no $filter_{file}$ (i.e., $sr = 0$); each single $search_{imp}$ and $search_{loc}$ strategy

does not extract too many import statements (mean 1.8) and keywords (mean 6.3). However, some extreme cases with 642 changed methods, 79 extracted keywords, and 28 different external import statements remain and require additional analyses to reduce the amount of data. For example, one can perform the API misuse detection approach only on the most suspicious methods (e.g., very complex methods, frequently changed methods), which are indicated by properties found in the change-based error detection domain (cf. Sect. 5.1).

3.3 Filtering analysis (RQ₂)

Methodology We conducted the analysis of all search and filter strategies (i.e., $search_{loc}$, $search_{imp}$, $filter_{file}$, and $filter_{method}$) as described in Sect. 2.2. As illustrated in Table 2, the analyzed strategies comprise 40 different configurations, all of which were evaluated for each of the 37 misuses from MUBench. This sums up to 1,480 different configurations. We implemented different scripts for conducting the strategies and obtained similar source files (i.e., $search_{loc}$) as described in Sect. 3.1. Regarding the $search_{imp}$ -step, we obtained the misused import statements by the meta-description of the misuses in the MUBench benchmark. With respect to the $filter_{file}$ strategy, we tested different values for the satisfaction ratio in the interval $[0, 1]$ and including the extremes $sr = 0.0$ (no keyword has to be matched) and $sr = 1.0$ (all keywords have to be matched). Since we estimate the best value for the satisfaction ratio without having too many configurations, we split the interval into four quarters, i.e., $[0, 0.25, 0.5, 0.75, 1]$. Finally, we applied the $filter_{method}$ -step by using the extracted sets of keywords.

For each configuration, we computed the *relative pattern frequency*. This describes how often a fixing pattern was found in the set of methods obtained from the retrieved source files. We obtained this by conducting the following five steps:

1. We manually distilled one or multiple variants of the fixing pattern using the known fix from the MUBench benchmark.
2. We generated the AUG for each fix
3. We generated the AUG of all methods obtained from the particular configuration.
4. We count how often the fixing pattern AUG is a subgraph in the set of AUGs obtained by a particular configuration.

Fig. 9 Distribution of the relative pattern frequency using different file search strategies grouped by API search strategy

5. We selected that pattern by the highest number of occurrences and divided that number by the number of AUGs.

Since the subgraph isomorphism problem is NP-hard, we only checked a relaxed condition. In particular, we only checked whether the set of nodes and the set of edges of the fixing pattern AUG is a subset of the set of nodes and the set of edges of the candidate AUG. Consequently, this introduces an overestimation, i.e., the real number of pattern occurrences might be lower. Therefore, we investigate the subsequent pattern mining ability in RQ₃.

We then compared the different configurations based on the relative pattern frequency. For that purpose, we applied the non-parametric Wilcoxon signed-rank test to determine whether the differences in the set of different configuration groups (if any exist) are significant. We chose this test instead of, for instance, the parametric t-test, since we cannot be sure that the relative pattern frequency follows a normal distribution. Particularly, the Wilcoxon signed-rank test has the null hypothesis that two paired groups originate from the same distribution. We reject the null hypothesis with $\alpha = 0.05$. For our analysis, these groups represent the relative pattern frequencies obtained from the strategies. Thus, the elements of the groups are paired by their misuse. Since the tests assume the elements of the single groups to be independent, we cannot simply split all frequency values from all configurations into two sets. For example, the 740 configurations using *filter_{method}* are not completely independent since 592 configurations used the same source files but with a different *filter_{file}* strategy. Therefore, to determine the real effect of a single strategy, we only compared those configurations using a single search strategy. When we determined the effect of one strategy the respective other filter strategies were left out. The concrete conditions under which the groups for comparison of the single strategies were obtained are depicted in Table 3.

Fig. 10 Distribution of the relative pattern frequency using different API search strategies grouped by each file search strategy

Table 4 Number of misuses per satisfaction ratio for which at least one fixing pattern was found

sr	Misuses with fixing patterns
0.00	22
0.25	21
0.50	21
0.75	18
1.00	4

Results In 748 out of 1480 cases, we obtained at least one similar source file fitting the criteria of the respective configuration. In 383 of those configurations, we found at least one occurrence of the fixing pattern. Regarding the misuses, for 33 misuses we found similar files. As denoted before, only in 31 cases, we were able to correctly extract the import of the misused API. Only for one of these 31 cases, we were not able to obtain any similar source files. For the 3 misuses, for which we found similar source files but not based on the misused API, our approach used other imports of shared third-party APIs. Consequently, it only selected source files that did not contain an occurrence of the fixing pattern. For 22 of those 30 misuses for which the misused API import was correctly extracted and for which we found similar source files, we found at least one fixing pattern with one of the 40 configurations.

Considering the strategies in detail, we were interested in which one has a significant positive impact in increasing the relative pattern frequency. Therefore, we first checked for how many misuses a particular strategy found at least one fixing pattern and second, whether the differences between single independent groups (denoted by Table 3) were significant w.r.t. the previously described test.

Fig. 11 Mean values of the relative pattern frequency among different file filter strategies (satisfaction ratio) grouped by certain search strategies

The $search_{loc}$ strategy distinguish between *internal* and *external* search. The internal $search_{loc}$ found similar files containing at least one occurrence of the fixing pattern for seven misuses, while the external $search_{loc}$ found them for 22 misuses. Thus quantitatively, we found more fixing patterns externally. This matches the observations made in previous work (Amann 2018). However, considering the distribution of relative frequencies in case a fixing pattern was found (Fig. 9), we can see that the mean relative frequency (indicated by the “x”-mark) is higher for the internal $search_{loc}$. This is true across both API search strategies. Nevertheless, the differences in the means may arise only from outliers of internal $search_{loc}$, while most considered misuses result in a relative pattern frequency of zero. Using the Wilcoxon signed-rank test, we could also not determine a significant difference in the distributions of the two file search strategies. Therefore, it indicates that using both, internal and external $search_{loc}$ in a cascaded manner could be useful. For example, before searching externally it might be worth first searching within the project itself. This relates to the idea of the plastic surgery theorem from the automatic program repair domain (Le Goues et al. 2019).

Next, we evaluated whether prior knowledge of the misused API imports has a significant positive impact on finding the fixing pattern. This is represented by the $search_{imp}$ -strategy. As our results show, we found fixing patterns for 22 misuses when using only the misused-imports- $search_{imp}$, compared to 17 misuses when using all-imports- $search_{imp}$. When considering the relative pattern frequency in Fig. 10, we cannot observe a significant difference between the two strategies using either internal or external code search. The Wilcoxon test also supported this observation. This indicates that prior knowledge of the misused API has only a moderate impact on increasing the number of fixing patterns. Considering that likely there exists no perfect method for identifying the misused API, it is reassuring to see that the results without such a method are not that much worth it.

(a) Internal/All Imports

(b) Internal/Misused Imports

Fig. 12 Distribution of the relative pattern frequency using different file filter strategies grouped by each API search strategy of internally found source code

$filter_{file}$ was applied for five different sr values (cf. Table 4). We observed that the number of misuses, for which we found at least one fixing pattern, was relatively stable (slightly drops from 22 to 18 misuses) with increasing sr , however, it drastically drops to four misuses for $sr = 1$.

The relative pattern frequency is usually constant (w.r.t. to the means) while dropping for $sr = 1$ (cf. Fig. 11). When comparing the distributions of srs combined with different strategies for $search_{loc}$ and $search_{imp}$, we determined a significant difference between $sr = 1$ and every other group of $search_{loc}$ and $search_{imp}$ except for applying internal $search_{loc}$ with all imports $search_{imp}$. Note that for the

(a) External/All Imports

(b) External/Misused Imports

Fig. 13 Distribution of the relative pattern frequency using different file filter strategies grouped by each API search strategy of externally found source code

internal $search_{loc}$ there are very few results so that the statistical tests may not be as reliable as for the external $search_{loc}$. Regarding the aspect of its lower relative frequency (e.g., in means), we can conclude that $filter_{file}$ with $sr = 1$ has a negative effect. While indicated by the mean values, the Wilcoxon test could not determine a significant difference in the distributions for $filter_{file}$ with $sr = 0.75$ and all distributions with a lower sr . For configurations using external $search_{loc}$, we could determine a significant difference between $filter_{file}$ with $sr = 0$ and $sr = 0.25$

(a) Internal File Search with an adapted y-axis for relative pattern frequency

(b) External File Search

Fig. 14 Distribution of the relative pattern frequency applying the method filter strategy grouped by each file and API search strategy

as well as $sr = 0$ and $sr = 0.5$. With respect to the mean values of these groups $filter_{file}$ has a slightly positive effect on the relative pattern frequency. Note that the respective median values are almost always zero for the different independent groups of $filter_{file}$. Therefore, we conclude that $filter_{file}$ usually has a moderate positive effect on the relative pattern frequency up to $sr = 0.5$ (cf. Figs. 12 and 13).

Finally, we analyzed the $filter_{method}$ strategy. Our findings are that by using $filter_{method}$ we could find fixing patterns for 21 misuses while finding fixes for 22

misuses when applying no $filter_{method}$. As depicted in Fig. 14, we observe a higher relative pattern frequency when applying the method filter strategy. Using the Wilcoxon signed-rank test, we could determine that the difference in distributions between applying and not applying $filter_{method}$ is significant. Note that due to the small number of unequal results for internal $search_{loc}$ the normal approximation used by the test might not hold and therefore should be taken with caution. However, this result indicates that $filter_{method}$ has a positive effect on increasing the relative pattern frequency.

We also analyzed the sr -values for each method that contains at least one keyword based on the raw source files (externally and internally collected as described in Sect. 3.1). The average sr for internally found methods is ≈ 0.11 while for externally found ones it is slightly higher with ≈ 0.19 . This supports our claim that the optimal sr for method filtering is lower. Thus, assuming that the average number of extracted keywords is 6.3 (cf. results of RQ₁), we consider a single keyword match is sufficient for method filtering since otherwise, we would remove too many methods (i.e., more than average). A more detailed analysis of the sr in the methods can be found in our replication package¹.

Regarding all 40 configurations, the overall best configuration (with mean ≈ 0.058) consisted of (1) using the internal $search_{loc}$, (2) using only the misused imports $search_{imp}$, (3) applying $filter_{file}$ with $sr = 0.25$ and (4) applying the $filter_{method}$. The best configuration using the external $search_{loc}$ (mean ≈ 0.026) used the misused imports $search_{imp}$, $filter_{file}$ with $sr = 0.75$, and applying $filter_{method}$.

Once again, we did not precisely determine the execution times for the single filtering steps. Based on the timestamps of the generated files, we can recall that the filtering for a single misuse takes at most two minutes. Note that this time excludes the searching and downloading of similar source files on Searchcode.

Implications Our findings show that even though an internal $search_{loc}$ might not find a pattern as often as the external $search_{loc}$, their distributions do not significantly differ. Therefore, we suggest a cascaded approach, which first searches for a pattern within the project and then in foreign projects.

Moreover, we conclude that prior knowledge of the misused API is likely to have only little benefits compared to the effort associated with the required preliminary analysis (cf. the results when applying $search_{imp}$). A possible explanation of why multiple imports still find fixing patterns is that these imports describe a kind of context, i.e., APIs that are frequently used together. Including this context representation in the search increases the chance of finding fixing patterns that fit the actual misuse.

When filtering source code, we found that $filter_{file}$ up to $sr = 0.5$ has a moderate effect on the relative pattern frequency, and thus may be applied to further reduce the number of source files.

In contrast, the $filter_{method}$ proved to be far more effective at increasing the relative pattern frequency. A possible explanation is that searching keywords in the method scope yields a more accurate representation of the misuse context than searching on the file level.

3.4 Mining analysis (RQ₃)

Methodology As versatile as the intermediate data representations are, there are also numerous ways of mining patterns from them. Usually, frequent pattern mining applies a variant of the well-known Apriori algorithm (Agrawal et al. 1993) with the extension of using closed pattern (Pasquier et al. 1999). A closed pattern is a pattern for which all super-patterns have lower support values. Mining algorithms differently rank the patterns according to some metrics, e.g. by support (how frequently does the pattern occur in the data set) or by confidence (how frequently do the elements of the pattern co-occur with the pattern). It was also found out that other metrics could be more effective in pattern mining (Le and Lo 2015).

In our experiments, we applied the API usage pattern miner by Amann, since it works with the data structure of API usage graphs (Amann 2018). This algorithm applies the Apriori algorithm with closed patterns on subgraphs by starting with individual AUG nodes and successively extending these depending on their neighbors in the AUG. For the extension process, the algorithm clusters isomorphic extensions, i.e., subgraphs. Note that, to cope with the graph isomorphism problem, they are using a graph vectorization heuristic. If the hash values of two graph vectors are equal, the graphs are considered isomorphic. Then a support threshold is used to identify which recurring sub-graphs should be reported as patterns. Further details can be found in the work by Amann (2018); Amann et al. (2019).

Further, we used the miner's *cross-method* support definition which counts only in how many different methods a pattern occurs.

In a first experiment, we only considered those 22 misuses from MUBench for which we could find at least one fixing pattern in the previous evaluation. Essentially, we wanted to assess the effect of our selected filter strategies compared to a mining process that does no pre-processing of the source files at all. For that purpose, we used the results of the source file collection as described in Sect. 3.1 as input for the un-filtered mining. Note that we conducted individual mining processes for internally and externally collected source files. For the filtered case, based on our previous analysis, we selected the filter strategy by searching all API imports, setting $sr = 0.5$, and applying the method filtering for both internally and externally collected source files.

Since we rank patterns regarding their support value, we have to set a minimum threshold. Based on our previous observation, we set the relative minimum support value (i.e., the ratio of absolute support to the number of all methods) for internal mining in both configurations (i.e., non-filtered vs. filtered) to 0.08. This value is selected based on the lower quartile of the distribution of configurations using internal file filtering which has a relative pattern frequency greater than zero and rounded this value down. For the external mining, we applied the same strategy based on the external filtering results and set the minimum relative support value to 0.004 for both mining configurations (i.e., non-filtered vs. filtered).

After mining, we sorted the patterns based on their absolute support (i.e., number of occurrences) and selected all patterns up to rank 20. Note that if multiple patterns had the same support, they share the same rank, while the next pattern with lower support has the next rank value increased by the number of patterns sharing

Table 5 Number of fixing patterns found in the Top@k patterns by mining without any filtering in the MUBench dataset

Misuse		Pattern in Top				add. is sub-pattern or equiv. pattern in Top			
		@ 1	@ 5	@ 10	@ 20	@ 1	@ 5	@ 10	@ 20
$\Sigma = 22$		3	3	4	8	3	6	7	8
1	alibaba_druid_1	–	–	–	–	–	–	–	–
2	alibaba_druid_2	–	–	–	–	–	–	–	–
3	android_rcs_rcsjta_1	–	–	–	–	–	–	–	–
4	apache_gora_56_1	–	–	–	–	–	–	–	–
5	apache_gora_56_2	–	–	–	–	–	–	–	–
6	bcel_101	–	–	–	–	–	–	–	–
7	jodatetime_269	–	–	E	E	–	–	E	E
8	jodatetime_361	E	E	E	E	E	E	E	E
9	jodatetime_362	E	E	E	E	E	E	E	E
10	jodatetime_363	E	E	E	E	E	E	E	E
11	logblock_2_15	–	–	–	–	–	–	–	–
12	mqtt_389	–	–	–	–	–	–	–	–
13	mqtt_390	–	–	–	–	–	–	–	–
14	tbuktu_ntru_473	–	–	–	–	–	–	–	–
15	tbuktu_ntru_474	–	–	–	–	–	–	–	–
16	tbuktu_ntru_475	–	–	–	–	–	–	–	–
17	testng_16	–	–	–	–	–	–	–	–
18	thebluealliancea_1	–	–	–	–	–	–	–	–
19	thomas_s_b_visualee_29	–	–	–	E	–	E	E	E
20	thomas_s_b_visualee_30	–	–	–	E	–	E	E	E
21	thomas_s_b_visualee_32	–	–	–	E	–	E	E	E
22	ushahidia_1	–	–	–	E	–	–	–	E

E pattern found in external source files;

I pattern found in internal source files;

I/E pattern found in internal and external source files

the previous rank. For example, if two patterns share rank 1 then the next pattern has rank 3.

Then, the first two authors independently reviewed the patterns and validated whether they depicted fixing ones. Particularly, we compared them with the fixing patterns that we already used in Sect. 3.3 and reviewed - if necessary - the respective documentation of the API to check for equivalent patterns. We further distinguished between the fixing type of *is pattern* and *is sub-pattern or equivalent pattern*. The rationale is that we rarely saw the pure form of the pattern by itself but rather the pattern being used as a sub-pattern in a bigger context. This is a result of closed pattern mining. Moreover, we found that some patterns to be semantically equivalent, which we did not consider as fixes in the first place. Note that if we classified the pattern

Table 6 Number of fixing patterns found in the Top@k patterns by mining with our predefined configurations in the MUBench dataset

Misuse		Pattern in Top				add. is sub-pattern or equiv. pattern in Top			
		@ 1	@ 5	@ 10	@ 20	@ 1	@ 5	@ 10	@ 20
$\Sigma = 22$		4	7	7	8	5	9	10	13
1	alibaba_druid_1	–	–	–	–	–	–	–	–
2	alibaba_druid_2	–	–	–	–	E	E	E	E
3	android_rcs_rcsjta_1	–	–	–	–	–	–	E	E
4	apache_gora_56_1	–	–	–	–	–	–	–	–
5	apache_gora_56_2	–	–	–	–	–	–	–	–
6	bcel_101	–	–	–	–	–	–	–	–
7	jodatetime_269	–	–	–	E	–	–	–	E
8	jodatetime_361	E	E	E	E	E	E	E	E
9	jodatetime_362	–	E	E	E	–	E	E	E
10	jodatetime_363	E	E	E	E	E	E	E	E
11	logblock_2_15	–	–	–	–	–	–	–	–
12	mqtt_389	–	I	I	I/E	–	I	I	I/E
13	mqtt_390	–	I	I	I	–	I	I	I
14	tbuktu_ntru_473	–	–	–	–	–	–	–	E
15	tbuktu_ntru_474	–	–	–	–	–	–	–	E
16	tbuktu_ntru_475	–	–	–	–	–	–	–	–
17	testng_16	–	–	–	–	–	–	–	–
18	thebluealliancea_1	–	–	–	–	–	–	–	–
19	thomas_s_b_visualee_29	I	I	I	I	I/E	I/E	I/E	I/E
20	thomas_s_b_visualee_30	I	I	I	I	I/E	I/E	I/E	I/E
21	thomas_s_b_visualee_32	–	–	–	–	–	–	–	–
22	ushahidia_1	–	–	–	–	–	E	E	E

E pattern found in external source files;

I pattern found in internal source files;

I/E pattern found in internal and external source files

as *is pattern*, we also marked it as *is sub-pattern or equivalent pattern*. Hence the right-hand side of Tables 5 and 6 can be considered as accumulated results.

We measured the agreement of the reviewers using Cohen's κ Cohen (1960). Reviewing the filtered results of internal code, we found no disagreements (i.e., $\kappa = 1$) regarding both cases, *is pattern* and *is sub-pattern or equivalent pattern*. For externally obtained patterns in the filtered case, we achieved a moderate agreement (i.e., $\kappa \approx 0.43$) for the *is pattern* review and substantial agreement (i.e., $\kappa \approx 0.69$) for the *is sub-pattern or equivalent pattern* review. Regarding the review of the unfiltered internal patterns, we could not compute κ since we only obtained five negative results among which the reviewers agreed. For external patterns the reviewers had perfect agreement (i.e., $\kappa = 1$) for the *is pattern* and almost perfect agreement (i.e.,

$\kappa \approx 0.84$) for *is sub-pattern or equivalent pattern*. Note that the reviews of the filtered and non-filtered results were done by the same reviewers but the latter one several months later, which could constitute a bias. During the second review process, we noticed that we falsely classified one pattern type in the filtered results. Thus, we reevaluated these patterns individually in the respective misuses and corrected our results. In our replication package, we explicitly marked these corrected results and our Cohen's κ computation also covers disagreements made in this reevaluation. Finally, we then discussed the conflicting points to understand the reasons why the respective reviewer accepted or rejected a particular pattern as a fix. Based on our discussion, we then agreed on a final classification.

In a second experiment, we analyze the effect of the filtering on each API usage in the AU500 dataset (Kang and Lo 2021). For that purpose, we conducted the internal and external code search for each method (including the exclusion steps as discussed before), applied the selected search and filter strategy (i.e., all imports, $sr = 0.5$, and method filtering), and mined usage patterns for the filtered and non-filtered case. Since a manual validation of patterns for those 500 API usages was infeasible, we applied the automatic violation detection from MUDetect by Amann (2018); Amann et al. (2019) to distinguish misuses from correct usages. While MUDetect employs many different violation techniques, we used a rather simple one that mines patterns with a certain minimum support and ranks the violations by the overlap between pattern and the usage. Particularly, this overlap between a pattern AUG p and an arbitrary API usage AUG u describes:

$$\text{overlap}(p, u) = \frac{|matchedNodes(p, u)| + |matchedEdges(p, u)|}{|nodes(p)| + |\{e \in edges(p) \wedge \text{notConnectedToAny}(e, nodes(p) \setminus nodes(u))\}|}$$

where *matchedNodes* and *matchedEdges* denote the set of matched nodes and edges between the two AUGs, *nodes* and *edges* denote the set of nodes and edges of an AUG, and *notConnectedToAny* is a predicate determining whether an edge e is not connected to any node in the set represented by the second parameter (i.e., $nodes(p) \setminus nodes(u)$). Particularly, it checks whether the edge e is not connected to any node from pattern p missing in the API usage AUG u . This way, we obtain an *overlap*¹² ranging from 0 to 1, where 0 is associated with no overlap while 1 depicts a perfect overlap. In our analysis, we marked all results with an overlap of exactly 0 or 1 as correct usages while usages with an overlap between 0 and 1 are denoted as misuses. We further configured MUDetect to find patterns with a minimal absolute support of 2 (for internally found code) and 10 (for externally found code), respectively. Note that in contrast to the previous experiments, we do not use relative support values, since MUDetect only supports absolute values, and thus, we selected those values based on typical absolute support values, we have seen in the previous experiments. To keep our experiment in a reasonable time frame, we timeout the external mining for each entry after ten minutes and the internal mining after five minutes. Since the labels of AU500 represent a ground truth of correct usages and misuses, we report the difference between the filtered and non-filtered case by the precision and the recall. To this end, we mark a detected misuse (i.e., an overlap

¹² In the MUDetect implementation this value is denoted as confidence. Since this can be confused with the confidence definition from the pattern mining domain, we denote this value as *overlap*.

Table 7 Results of the misuse detection on the AU500 dataset using the violation detection technique from MUDetect with number of true positives (#tp), false positives (#fp), true negatives (#tn), false negatives (#fn), precision, and recall

Applied filtering	<i>search_{loc}</i>	#tp	#fp	#tn	#fn	Precision (%)	Recall (%)
No	External	8	16	369	107	33.33	6.96
	Internal	1	10	375	114	9.09	0.87
	Both	9	26	359	106	25.71	7.83
Yes	External	8	12	373	107	40.0	6.96
	Internal	9	27	358	106	25.0	7.83
	Both	13	31	354	102	29.55	11.3

between 0 and 1) as true-positive if it was also labeled as misuse in the ground truth and as false-positive if not. In case the detection did not determine a misuse, we mark it as true-negative if it was labeled as a correct usage and as false-negative otherwise.

Results Tables 5 and 6 depict the results of our evaluation of the MUBench dataset regarding the API usage pattern mining for the non-filtered and the filtered mining configurations, respectively.

In the unfiltered configuration, the miner found the fixing pattern in its ‘pure’ form for four misuses in the Top@10 and eight misuses in the Top@20 most frequent patterns, respectively. If we relax the requirement so that the pattern can be a sub-pattern or an equivalent (sub-)pattern, this number increases to seven for Top@10. For the Top@20 case, it remains constant.

In the filtered configuration the ‘pure’ form was found for seven misuses in the Top@10 and eight misuses in the Top@20. Concerning the relaxed case, the numbers increase to 10 and 13 misuses for the Top@10 and Top@20 frequent patterns, respectively.

Since our mining used the same original data and applied the same configuration (i.e., minimum support), we observed that our selected searching and filtering strategy had a positive effect on the number of detected fixing patterns for the analyzed misuses. We analyzed whether the difference in the number of occurrences of fixing patterns in the Top@20 is significant. For that purpose, we used the χ^2 -test ($\alpha = 0.05$), which due to the small sample size of 22 analyzed misuses, was corrected with the Yates correction (Yates 1934). Based on this test the difference is not significant.

However qualitatively, we observed in the unfiltered case that fixing patterns were solely inferred from the externally found source code, while in the filtered case it found patterns from the internal code as well. We assume that this is due to the fact that the external code is being filtered to some degree (i.e., it was collected by searching with related imports on Searchcode), while the internal code was simply used as it is. Moreover, the unfiltered mining only found fixing patterns for misuses from three different projects, while the filtered mining does so for seven. Thus, this

indicates that the positive effect of our strategy is more general in terms of different projects.

In more detail, the filtered case found more patterns, sub-patterns, or equivalent patterns in external code than in internal code, i.e., 12 compared to 4. For three misuses (i.e., `mqtt_389`, `thomas_s_b_visualee_29`, and `thomas_s_b_visualee_30`) both strategies, i.e., internal and external, found a fixing pattern. This indicates that the combination of the two configurations, namely, internal and external code search, is beneficial.

Nevertheless, for nine misuses, the filtered case did not find a fixing pattern in the Top@20 highest-ranked patterns. Of these, the miner could not obtain any patterns for three misuses, namely, `apache_gora_56_2`, `testng_16`, and `thomas_s_b_visualee_32`. The miner failed to find patterns for the external configuration only in these three cases, while for the internal configuration, it did not find any patterns for 17 misuses. For seven misuses, we found patterns but could not retrieve a fixing pattern in the set of Top@20 ranked patterns.

For the cases `jodatetime_269`, `thomas_s_b_visualee_29-32`, and `ushahidia_1` the filtering actually removed too many valid occurrences so that either the fixing patterns were not ranked as highly as without filtering or did not find the patterns at all.

During the analysis of the filtered results, we made the observation that many patterns tend to be very close to the fixing pattern but miss some essential parts to be considered as a solution. Moreover, we observed many very similar patterns distributed among the highest-ranked patterns. By merging these sub-patterns one may re-build the fixing pattern.

Regarding the execution time of the miner step, we can roughly recall from the timestamps of the generated files that a single execution in most cases takes at most a minute regardless of whether the results were filtered or not. Only in a single case (filtered external source files of `thomas_s_b_visualee_29`), the mining took almost 18 minutes.

Finally, we analyzed the effect of the previous searching and filtering on the violation detection by MUDetect (Amann 2018; Amann et al. 2019) applied on the AU500 dataset by Kang and Lo (2021). Table 7 summarizes our results, by individually depicting the results of externally and internally searched code as well as a cascaded approach by applying both search strategies. Please note, that the row showing the values for both strategies does not necessarily represent the sum of the external and internal row since the sets of detected misuses from external and internal search can overlap. Moreover, we observed that we could apply our approach only to 480 out of those 500 API usages in the dataset. In our computation, we considered these 20 entries as if the approach did not detect a misuse (i.e., they add either to #tn or #fn).

In general, we can observe a positive effect of the filtering. Particularly, the number of true positives for the internal search increased, and thus, their precision and their recall. When analyzing the reasons, we found out that this was an effect of the decreased number of AUGs to generate and mine from. In case we did not use the filtering strategy, the number of methods was too large, and thus, only 20 API usages could be analyzed using the internal, non-filtered case, while

the others were interrupted by the timeout. In contrast, for the filtered case, we were enabled to analyze 453 API usages. Regarding the external case, we could slightly improve the precision due to a lower number of false positives, while the recall was consistent. Interestingly, even though the number of true positives is equal for both external strategies, the detected API misuses differ. In detail, both strategies, the filtered and non-filtered external search, independently enabled the violation detection to correctly detect a misuse for a common set of five misuses. Every individual strategy, however, enabled the violation detection to find misuses for three API usages, which were not detected using the respective other strategy. The cause of this is that the filtering decreases the support values for those three patterns too much pushing it below the minimum support. In the other, non-filtered case, we observed that the mining was interrupted by a timeout. For the cascaded approach, we observed a lower precision than for the external setup, while the recall is higher. When applying the filtering, we observe an increase for both precision (i.e., +3.84%) and recall (i.e., +3.47%) compared to the non-filtered case. A subsequent χ^2 -test ($\alpha = 0.05$) found the difference of the true positive values to be non-significant. While we found the precision to be comparable with previous results (Amann et al. 2019; Kang and Lo 2021) the recall is rather low. We will discuss the implications of this observation in the following paragraph.

Implications In both experiments on the MUBench (Amann et al. 2016) and the AU500 dataset (Kang and Lo 2021), we could not determine a significant effect on the difference in the number of found fixing patterns and true positives when applying filtering compared to no filtering. However, from a qualitative perspective, we found that filtering still has a positive effect, for instance, we found fixing patterns for more distinct projects and based on internally collected source files. Its main effect, however, is the reduction of code samples, allowing faster mining on fewer examples without negatively influencing the misuse detection capability.

Further, we found that our search strategy is not perfect. For MUBench, we found patterns for around 59% of the considered misuses and around 35% of all 37 misuses. For the AU500 dataset, the maximum achieved recall was 11.3%. Therefore, we will consider in future work whether other code search strategies perform better. For instance in the work by Kang and Lo (2021), the authors applied the recently published artifact AUSEarch Asyrofi et al. (2020) and achieved higher recall values. Moreover, we applied very simple ranking strategies (support and violation overlap). As it has been shown by previous work, other metrics could further improve the results (Le and Lo 2015; Amann 2018). Additionally, recent research has come up with potential further datasets for API misuses (Nielebock et al. 2021; Kechagia et al. 2021), which provide a more diverse set of validation data.

We could confirm that the combination of both code search strategies contributed to retrieving more fixing patterns, while the external strategy found patterns for more misuses than the internal one in the MUBench dataset. In the AU500 dataset, both had an almost equal number of true positives while the external search achieved usually a higher precision.

Additionally, we observed many very similar patterns as well as incomplete patterns in the ranking of the filtered results of MUBench. By clustering these patterns,

similar to previous work (Zhong et al. 2009), we could decrease the number of similar patterns to a lower number of clusters. The patterns within these clusters would then represent a set of different possible options for applying an API in a particular context. Depending on the strategy, one could simply pick the most frequent pattern from each cluster or merge patterns regarding some heuristic to re-combine the missing parts. This will be part of future work.

4 Threats to validity

The validity of our evaluation could be subject to different threats. With respect to related literature (Siegmond et al. 2015), we consider threats to *internal* and *external validity*.

Internal validity Internal validity describes to which degree we can trust our results. Particularly, errors made in our process could harm the robustness of our results.

In our concept, we rely on identifying similar source code which is likely to contain fixes for a particular misuse. For external search, we used *SearchCode*, which leverages data from different code repositories. Depending on the concrete misuse, significant time may have passed between the misuse introduction and our similar code search. Therefore, the discovered code may not yet have existed at the time the misuse was first introduced. This may imply that the changed-based approach might find fewer fixes when executed just-in-time.

Moreover, the search results are biased through *Searchcode*'s search algorithm, as a different external search might perform differently. However, we found the subsequent search and filter strategies to perform similarly among externally and internally found code sets. In fact, we found the same strategies for both sets as best fitting for the mining process of RQ₃. Additionally, we kept a consistent data set for all search strategies so that we assume this effect to be almost equal for all strategies. Based on this assumption, we still expect that our results express which strategy works best.

Even though we filtered the similar code to exclude any files originating from the source project of the misuse, our process cannot guarantee that we do not find code originating from forks of that project. However, these threats are to a large degree mitigated when, in particular for frequently used APIs, we find many similar usages.

Our process is capable of inferring fixing patterns. However, we can guarantee neither the patterns' correctness nor their completeness. Regarding the former, we did not check whether the fixing pattern would introduce new errors. The latter is due to the fact that we can never be sure to have found all possible variations of a fix. To some extent, this threat is again mitigated when the fixing patterns exhibit higher support values, as this would likely favor the relevant and more general fix variants.

We applied the relative pattern frequency to compare different search and filter strategies in their ability to find fixing patterns. However, this metric could be biased in case we retrieved a very low number of source files. For example,

the number of internally found files is usually lower than those of external ones. Thus, comparative statistics such as the differences in means might be only by chance.

Finally, the first two authors independently validated whether the mined patterns represented the respective fix or not. However, as manual validation always has subjective aspects, this may introduce bias or noise into our evaluation. Moreover, the two separate review phases were separated by a gap of several months, which may bias the agreement metrics since the reviewers had a learning experience from the first review. For this reason, we published all our results, data, and scripts as a replication package¹ and invite other researchers to re-validate our findings.

External validity External Validity describes to which degree our results generalize to unknown data, i.e. other misuses from different projects.

Our methodology only considers static and intra-procedural API usage patterns, i.e., patterns within the scope of a single method declaration. However, API usage patterns may also be scattered among several methods, i.e., inter-procedural or may only be detected when the code is explicitly executed (i.e., dynamic pattern inference). Since our methodology does not detect such inter-procedural or dynamically inferred patterns, our results currently only refer to static, intra-procedural API usage pattern miners.

Our MUBench-based case study features many similar misuses. Regarding this point, Sven Amann, one of the authors of MUBench, noted “*The benchmark dataset may not be representative for API misuses in the wild*” (Amann 2018, p. 75). Therefore, if MUBench is subject to this limitation, the same is likely true for our case study.

Moreover, most of our evaluation considers only a small set of 37 real API misuses from MUBench. While specifically preprocessed to reduce potential bias (e.g., removing duplicates), future work may validate our results on larger datasets of API misuses such similar to ours on the AU500 dataset (Kang and Lo 2021) as well as recently published ones (Nielebock et al. 2021; Kechagia et al. 2021).

Our method and analysis refer only to API misuses in Java. We did not study whether this method would perform similarly for other programming languages, however, if the keyword extraction is adapted, this is arguably the case for other procedural and object-oriented languages. Regarding other programming paradigms, an adaptation would likely also require conceptual adaptations.

5 Related work

Our work relates to four further software engineering fields, which we consider in the following.

5.1 Change-based error detection

The idea of detecting bugs based on commits is not new. Originally, these works investigated metrics indicating suspicious commits that introduced bugs. This is also known as just-in-time bug detection. Mockus and Weiss used change properties such as size or diffusion (e.g., number of distinct files that have to be changed) and build a model based on a logistic regression to estimate the probability of an error (Mockus and Weiss 2000). Sliwski et al. introduced the SZZ-algorithm by which they could identify bug-introducing commits by using the version control system (Śliwerski et al. 2005). We also used this algorithm to determine the API misuse-introducing commits. Kim et al. trained a support vector machine (SVM) based on information from the source code metadata and achieved an accuracy of 78% to detect bug-inducing commits (Kim et al. 2008). However, they found the model to be too project-specific to be globally usable. Instead Kamei et al. (2013), similarly to Mockus and Weiss' idea, used a logistic regression and found a generic model achieving an average accuracy of 68%. Their model indicates that commits with files that get large and frequent changes are associated with introducing bugs. An and Khomh (2015) confirmed this observation regarding the changes' sizes and found further correlations with low developers-experience, longer commit messages, and changes distributed across multiple files. Augmenting these general characteristics of suspicious commits, we have begun to investigate the influence of API-specific information as shown in our preliminary work (Nielebock et al. 2018). Other tools like ChangeLocator Wu et al. (2017) and ChangeRanker Guo et al. (2020) improve the accuracy by applying information from automatically collected crash reports. Since we aim to detect API misuses at the time of the commit, this data is usually not available. Other approaches detect bugs by modeling code changes as logic rules and try to detect exceptions from these rules within the code change (Kim and Notkin 2009). Similarly, some automated bug detection and repair techniques use previous or historical bug fixes (Kim et al. 2006; Sun et al. 2010; Le et al. 2016; Long and Rinard 2016). In contrast, our approach considers API usages at the time of a commit. However, we also worked on re-using previous fixes to detect misuses in other repositories (Nielebock et al. 2020a). Whether using historical fixes is beneficial for our approach will be the subject of our further work.

5.2 API selection and usage recommendation

API recommendation aims to assist developers in selecting suitable APIs for their use-case and also correctly applying those APIs. While in our case the API selection is fixed by the given source code, some of the search strategies from this field inspired our approach.

Saul et al. implemented the FRAN (Find with RANdom walks) algorithm, which, given an API and a particular function, finds closely related functions of the same API (Saul et al. 2007). The *Prospector* assists developers in creating objects by recommending code based on the desired type of the object and likely-relevant parameters (Mandelin et al. 2005). Similarly, Chan et al. use a subgraph-based

algorithm to find relevant code regarding a textual query (Chan et al. 2012). Other approaches use textual input from feature requests (Thung et al. 2013b) or search for textual queries using additional sources such as code documentations or Q&A web-pages (Lv et al. 2015; Rahman et al. 2016). Thung et al. also developed an approach that suggests complete libraries based on the respective APIs currently used in a client project (Thung et al. 2013a). The MUSE approach finds usage examples for individual specified API methods by means of static slicing for simplification and various heuristics for ranking (Moreno et al. 2015). In our approach, we include data types, called methods, and import statements into the context of the API usage and use these to find and filter similar API usages. Note that usually, we do not use all information at the same time. Moreover, we do not expect to have access to the declaration of those functions that are called by the query function. A very recent approach uses API embeddings (Nguyen et al. 2017; Chen et al. 2019), or joint natural text and API embeddings (Huang et al. 2018) to find similar API usages. These techniques enable finding semantically equivalent APIs even if they do not share syntactical similarities. This, however, requires the training of a neural network, which would not scale for our envisioned use case. Moreover, applying semantically equivalent APIs would require significant changes to the code (i.e., substituting complete libraries) and would thus potentially introduce further sources of bugs.

5.3 Code search and code recommendation

Code search is usually referred to as the task performed by developers to retrieve code of varying size from code snippets up to complete packages (Gallardo-Valencia and Sim 2009). The main motivations for code search are code reuse, -repair, -understanding, -location, impact analysis, or finding suitable third-party libraries (Gallardo-Valencia and Sim 2009; Sadowski et al. 2015; Xia et al. 2017). This indicates that the code search is heavily domain and use-case-specific. However, still, developers tend to prefer general-purpose search engines for code search (Sadowski et al. 2015; Rahman et al. 2018). In contrast, our goal was to find similar code examples without requiring human interaction.

Several automatic approaches aim to facilitate the search process and the quality of the results. These approaches range from using a domain-specific languages (Paul and Prakash 1994), the code context (Holmes and Murphy 2005; Sahavechaphan and Claypool 2006), test cases (Lemos et al. 2007), static and dynamic specifications in the form of method signatures and test cases (Reiss 2009), the documentation and slicing techniques (Kim et al. 2010), textual matching based on different ranking and natural language processing mechanism (McMillan et al. 2011), input/output code examples using SMT-solver (Satisfiability modulo theories solver) (Stolee et al. 2014), or learned neural code embeddings (Gu et al. 2018). All these approaches usually aim to find accurate search results. In contrast, our approach can cope with a certain degree of ‘noise’ in the search results due to subsequent filter and mining steps. This allows us to keep the search algorithm as lightweight as possible.

During the development of this work, the AUSEarch tool was published using user queries to find similar API usages (Asyrofi et al. 2020). Their technique uses

type resolutions to retrieve better matching samples. It was successfully applied by Kang and Lo (2021) when applying their API misuse detector. However, they applied the user queries manually, while our approach is intended to run fully automatically. In future work, we incorporate AUSEarch in the overall process.

Note that the keyword search using the context is very similar to the notions presented in Strathcona Holmes and Murphy (2005) and the XSnippet tool (Sahavechaphan and Claypool 2006). These tools, like our approach, used inheritance, type- and method-call information from the method declaration to retrieve similar code. Unfortunately, both tools were not available at the time we conducted our experiments.

Some recent work on automated program repair by Xin and Reiss used code search to reduce the typical huge search space for patches (Xin and Reiss 2017, 2019). In their work, they made the interesting observation that it is worth using different search strategies for internal and external code. We will consider this, as a potential extension of our work.

5.4 Code clone detection

Finding similar source code is related to retrieving syntactically and semantically equivalent source code, namely, code clones. The motivation for detecting code clones is manifold and includes reducing maintenance effort, detecting plagiarism, code compaction, analyzing software evolution, and bug detection. A number of tools were developed to support clone detection (Koschke 2007; Roy et al. 2009). Deckard is well-known and uses parse tree vectorization to find clones (Jiang et al. 2007). In the recent research other tools aim to find different levels of clones ranging from Type-3 (near-miss) to Type-4 (same functionality) (Sajnani et al. 2016; White et al. 2016; Saini et al. 2018). For our use-case, however, clone detection is less relevant, since the fixing code for a target API misuse by definition cannot be a clone of the former.

6 Conclusion and further work

Recent research came up with a variety of automatic API misuse detectors that rely on the idea of inferring correct API usages from existing code samples. However, we discovered that most approaches do not consider how these code samples can be collected in practice. Therefore, this paper introduces a new approach to collect and improve the input source code files for API usage pattern mining to more effectively find patterns for API misuses detection. This approach uses a program-change analysis combined with similar code search and a filtering strategy. We also introduce a concept with which this approach can easily be integrated into an ordinary continuous integration process. This concept has two advantages: First, it applies API misuse mining in a just-in-time manner whenever developers commit changes in their code to the version control system. Second, to improve the results of the

miner, based on the concrete change, it applies different search and filter strategies to increase the relative frequencies of true positive patterns in a set of code samples before mining.

In our experiments, we determined the overall best search and filter strategy by analyzing 37 well-known API misuses from the MUBench dataset (Amann et al. 2016) and selecting that strategy that achieved the highest relative frequency. Using this strategy, we analyzed the effect on the pattern mining and misuse detection using the tooling by Amann (2018); Amann et al. (2019) and two different API misuse datasets (Amann et al. 2016; Kang and Lo 2021).

Our main findings are:

1. Considering only changed methods that modified the usage of third-party libraries can effectively reduce the number of methods to investigate (i.e., on average reduction of 86.4% of methods to be analyzed).
2. Both, internal (i.e., within the project) and external (i.e., in other projects) code search contribute to more fixing patterns being found.
3. Including knowledge of which API was misused into the search has only a negligible effect and therefore it is generally sufficient to search for similar API usages without exactly knowing the misused API.
4. File filtering with keywords has only a moderate effect on increasing the relative frequency of patterns, while method filtering with keywords has a significant positive effect without removing too many real patterns.
5. In comparison to non-filtered results, with our strategy we retrieved more patterns that can fix a misuse. Even though the difference is not significant in a quantitative manner, we qualitatively observed that we found more patterns regarding different projects and for internally collected source code. Nevertheless, our search strategy is improvable, for instance, by recent work on similar API code search (Asyrofi et al. 2020).

Based on our results and existing related work, we plan the following additional steps to further integrate this knowledge into a full-fledged API misuse detection and correction tool.

Detecting API-misuse commits While we found that analyzing API usage on the commit level reduces the effort in terms of the number of methods to analyze, we still have some huge outliers. Moreover, related work (Mockus and Weiss 2000; Kim et al. 2008; Kamei et al. 2013; An and Khomh 2015) also suggests techniques for discriminating bug-containing from bug-free commits. We plan to include these techniques as a further pre-processing step so that API change analysis is only conducted when the commit is determined as suspicious, such as indicated by our prior work (Nielebock et al. 2018). Additionally, to cope with huge changes consisting of several hundreds of potentially misuse-containing methods, we could also transfer these techniques to the method scope, i.e., detecting particularly suspicious methods (e.g., by the frequency of method's changes).

Code search and filtering Up to now, we have aimed for a lightweight search and filter process. That means, we avoided costly static and dynamic code analyses. Assuming that we can further reduce the absolute number of methods to analyze, it could be worth introducing further analysis as those presented in Sect. 5.3, for instance, by applying AUSearch Asyrofi et al. (2020). We also made the observation that for some APIs there exist only a few code examples. To solve this problem a recent idea is to represent API usages in the form of a learned vector embedding, such as API2Vec Nguyen et al. (2017). This embedding depicts semantic relations between API usages as vector operations, such as $v(ListIterator.hasNext) \approx v(StringTokenizer.hasMoreTokens) - v(StringTokenizer.nextTokens) + v(ListIterator.next)$. A first notion of how this can be leveraged to map verified usage patterns from well-known APIs to equivalent, less-known APIs is published in Nielebock et al. (2020b).

API misuse detection Our analysis shows that in case we have a misuse introducing commit and we specifically investigate the misuse-containing method, we can improve the relative frequency of true positive patterns in retrieved similar API usages. However, we also have to consider to what degree our approach may falsely classify correct API usages as misuses. Our analysis on the AU500 dataset indicates that further steps for increasing the process precision are necessary. Finally, we need a human-based study on the usability of our approach, since there are a number of factors that may hamper developer's acceptance.

API misuse repair We envision a full-fledged process that not only detects API misuses automatically but also suggests patches. For that purpose, we have to incorporate the fixing pattern (i.e., that detected the misuse) back into the original code. This requires further post-processing such as mapping variables. Moreover, the patches need to be validated, which other automated program repair approaches typically do by using test suites. However, we do not have such a test suite or cannot ensure whether these tests check the API misuse behavior (Le Goues et al. 2019). Thus, we require human intervention. To minimize that effort, we have to ensure that we have only a few patch candidates with a high rate of true positives.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, 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 licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence 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. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agrawal, A., Menzies, T.: Is "Better Data" Better Than "Better Data Miners"? On the Benefits of Tuning SMOTE for Defect Prediction. In: Proceedings of the 40th International Conference on Software Engineering (ICSE), ACM, New York, NY, USA, pp. 1050–1061 (2018). <https://doi.org/10.1145/3180155.3180197>
- Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. *SIGMOD Rec.* **22**(2), 207–216 (1993). <https://doi.org/10.1145/170036.170072>
- Allamanis, M., Sutton, C.: Mining Idioms from Source Code. In: Proceedings of the 22nd International Symposium on the Foundation of Software Engineering (FSE), ACM, New York, NY, USA, pp. 472–483 (2014). <https://doi.org/10.1145/2635868.2635901>
- Amann, S.: A Systematic Approach to Benchmark and Improve Automated Static Detection of Java-API Misuses. PhD thesis, Technische Universität Darmstadt (2018). URL <http://tubiblio.ulb.tu-darms.tadt.de/106302/>
- Amann, S., Nadi, S., Nguyen, H.A., Nguyen, T.N., Mezini, M.: MUBench: A Benchmark for API-misuse Detectors. In: Proceedings of the 13th International Working in Mining Software Repositories (MSR), ACM, New York, NY, USA, pp. 464–467 (2016). <https://doi.org/10.1145/2901739.2903506>
- Amann, S., Nguyen, H.A., Nadi, S., Nguyen, T.N., Mezini, M.: A systematic evaluation of static API-misuse detectors. *IEEE Trans. Softw. Eng.* (2018). <https://doi.org/10.1109/TSE.2018.2827384>
- Amann, S., Nguyen, H.A., Nadi, S., Nguyen, T.N., Mezini, M.: Investigating Next Steps in Static API-Misuse Detection. In: Proceedings of the 16th International Working on Mining Software Repositories (MSR), IEEE Press, pp. 265–275 (2019). <https://doi.org/10.1109/MSR.2019.00053>
- Ammons, G., Bodík, R., Larus, J.R.: Mining specifications. In: Proceedings of the 29th Symposium Principles of Programming Language (POPL), ACM, New York, NY, USA, pp. 4–16 (2002). <https://doi.org/10.1145/503272.503275>
- An, L., Khomh, F.: An Empirical Study of Crash-inducing Commits in Mozilla Firefox. In: Proceedings of the 11th International Conference on Predictive Models in Software Engineering (Promise), ACM, New York, NY, USA, pp. 5:1–5:10 (2015). <https://doi.org/10.1145/2810146.2810152>
- Asyofi, M.H., Thung, F., Lo, D., Jiang, L.: AUSEarch: Accurate API Usage Search in GitHub Repositories with Type Resolution. In: Proceeding 27th International Conference on Software Analytic Evolution, Reeng. (SANER), pp. 637–641 (2020). <https://doi.org/10.1109/SANER48275.2020.9054809>
- Chan, W.K., Cheng, H., Lo, D.: Searching Connected API Subgraph via Text Phrases. In: Proceedings 20th International Symposium Found Software Engineering (FSE), ACM, p. 10 (2012). <https://doi.org/10.1145/2393596.2393606>
- Chen, C., Xing, Z., Liu, Y., Ong, K.L.X.: Mining Likely Analogical APIs across Third-Party Libraries via Large-Scale Unsupervised API Semantics Embedding. *IEEE Transactions Software Engineering*, pp. 1–1 (2019). <https://doi.org/10.1109/TSE.2019.2896123>
- Cohen, J.: A coefficient of agreement for nominal scales. *Educ. Psychol. Meas.* **20**(1), 37–46 (1960). <https://doi.org/10.1177/001316446002000104>
- Dyer, R., Nguyen, H.A., Rajan, H., Nguyen, T.N.: Boa: A Language and Infrastructure for Analyzing Ultra-large-scale Software Repositories. In: Proceedings 35th International Conference Software Engineering (ICSE), IEEE Press, Piscataway, NJ, USA, pp. 422–431 (2013). <https://doi.org/10.1109/ICSE.2013.6606588>
- Ernst, M.D., Cockrell, J., Griswold, W.G., Notkin, D.: Dynamically discovering likely program invariants to support program evolution. *IEEE Trans. Softw. Eng.* **27**(2), 99–123 (2001). <https://doi.org/10.1109/32.908957>
- Ernst, M.D., Perkins, J.H., Guo, P.J., McCamant, S., Pacheco, C., Tschantz, M.S., Xiao, C.: The daikon system for dynamic detection of likely invariants. *Sci. Comput. Progr.* **69**(1), 35–45 (2007). <https://doi.org/10.1016/j.scico.2007.01.015>, Special Issue on Experimental Software and Toolkits
- Frolin S Ocariza, J., Bajaj, K., Pattabiraman, K., Mesbah, A.: An Empirical Study of Client-Side JavaScript Bugs. In: Proceeding 7th International Symposium on Empirical Software Engineering and Measurement (ESEM), pp. 55–64 (2013). <https://doi.org/10.1109/ESEM.2013.18>
- Gabel, M., Su, Z.: Javert: Fully Automatic Mining of General Temporal Properties from Dynamic Traces. In: Proceedings 16th International Symposium Found Software Engineering (FSE), ACM, pp. 339–349 (2008). <https://doi.org/10.1145/1453101.1453150>

- Gallardo-Valencia, R.E., Sim, S.E.: Internet-Scale Code Search. In: Proceedings 1st Work. Search-Driven Dev.-Users, Infrastruct., Tools Evaluation (SUITE), IEEE, pp. 49–52 (2009). <https://doi.org/10.1109/SUITE.2009.5070022>
- Gousios, G.: The GHTorrent Dataset and Tool Suite. In: Proceedings 10th International Work Min Software Repository (MSR), pp. 233–236 (2013). <https://doi.org/10.1109/MSR.2013.6624034>
- Gu, X., Zhang, H., Kim, S.: Deep Code Search. In: Proceeding 40th International Conference Software Engineering (ICSE), IEEE, pp. 933–944 (2018). <https://doi.org/10.1145/3180155.3180167>
- Guo, Z., Li, Y., Ma, W., Zhou, Y., Lu, H., Chen, L., Xu, B.: Boosting crash-inducing change localization with rank-performance-based feature subset selection. *Empir. Softw. Eng.* pp. 1–46 (2020). <https://doi.org/10.1007/s10664-020-09802-1>
- Holmes, R., Murphy, G.C.: Using Structural Context to Recommend Source Code Examples. In: Proceedings 27th International Conference Software Engineering (ICSE), ACM, New York, NY, USA, pp. 117–125 (2005). <https://doi.org/10.1145/1062455.1062491>
- Hou, D., Li, L.: Obstacles in Using Frameworks and APIs: An Exploratory Study of Programmers' Newsgroup Discussions. In: Proceedings 19th International Conference Program Comprehension. (ICPC), IEEE, pp. 91–100 (2011). <https://doi.org/10.1109/ICPC.2011.21>
- Huang, Q., Xia, X., Xing, Z., Lo, D., Wang, X.: API Method Recommendation without Worrying about the Task-API Knowledge Gap. In: Proceedings 33rd International Conference Automatic Software Engineering (ASE), ACM, New York, NY, USA, pp. 293–304 (2018). <https://doi.org/10.1145/3238147.3238191>
- Jiang, L., Mishnerghi, G., Su, Z., Glondu, S.: DECKARD: Scalable and Accurate Tree-Based Detection of Code Clones. In: Proceedings 29th International Conference Software Engineering (ICSE), IEEE Computer Society, pp. 96–105 (2007). <https://doi.org/10.1109/ICSE.2007.30>
- Kamei, Y., Shihab, E., Adams, B., Hassan, A.E., Mockus, A., Sinha, A., Ubayashi, N.: A large-scale empirical study of just-in-time quality assurance. *IEEE Trans. Softw. Eng.* **39**(6), 757–773 (2013). <https://doi.org/10.1109/TSE.2012.70>
- Kang, H.J., Lo, D.: Active learning of discriminative subgraph patterns for API misuse detection. *IEEE Trans. Softw. Eng.* (2021). <https://doi.org/10.1109/TSE.2021.3069978>
- Kechagia, M., Mechtaev, S., Sarro, F., Harman, M.: Evaluating automatic program repair capabilities to repair API misuses. *IEEE Trans. Softw. Eng.* (2021). <https://doi.org/10.1109/TSE.2021.3067156>
- Kim, J., Lee, S., Hwang, S.W., Kim, S.: Towards an Intelligent Code Search Engine. In: Proceedings 24th AAAI Conference Artificial Intelligence (AAAI) (2010). URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/download/1691/2212>
- Kim, M., Notkin, D.: Discovering and Representing Systematic Code Changes. In: Proceedings 31st International Conference Software Engineering (ICSE), IEEE Computer Society, USA, pp. 309–319 (2009). <https://doi.org/10.1109/ICSE.2009.5070531>
- Kim, S., Pan, K., Whitehead, E.E.J.: Memories of Bug Fixes. In: Proceedings 14th International Symposium Found Software Engineering (FSE), ACM, New York, NY, USA, pp. 35–45 (2006). <https://doi.org/10.1145/1181775.1181781>
- Kim, S., E James Whitehead, J., Zhang, Y.: Classifying software changes: clean or buggy? *IEEE Trans. Softw. Eng.* **34**(2), 181–196 (2008). <https://doi.org/10.1109/TSE.2007.70773>
- Koschke, R.: Survey of Research on Software Clones. In: Koschke, R., Merlo, E., Walenstein, A. (eds.) Dagstuhl Semin. Proc.: Duplic., Redundancy, and Similarity Softw., Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, Dagstuhl, Germany, no. 06301 in Dagstuhl Seminar Proceedings (2007), URL <http://drops.dagstuhl.de/opus/volltexte/2007/962>
- Le, T.D.B., Lo, D.: Beyond Support and Confidence: Exploring Interestingness Measures for Rule-Based Specification Mining. In: Proceedings 22nd International Conference Software Analytical Evolution, Reeng. (SANER), pp. 331–340 (2015). <https://doi.org/10.1109/SANER.2015.7081843>
- Le, X.B.D., Lo, D., Le Goues, C.: History Driven Program Repair. In: Proceedings 23rd International Conference Software Analytical Evolution Reeng. (SANER), vol 1, pp. 213–224 (2016). <https://doi.org/10.1109/SANER.2016.76>
- Le Goues, C., Weimer, W.: Measuring code quality to improve specification mining. *IEEE Trans. Softw. Eng.* **38**(1), 175–190 (2012). <https://doi.org/10.1109/TSE.2011.5>
- Le Goues, C., Pradel, M., Roychoudhury, A.: Automated program repair. *Commun. ACM* **62**(12), 56–65 (2019). <https://doi.org/10.1145/3318162>
- Lemos, O.A.L., Bajracharya, S.K., Ossher, J., Morla, R.S., Masiero, P.C., Baldi, P., Lopes, C.V.: CodeG-enie: Using Test-Cases to Search and Reuse Source Code. In: Proceedings 22nd International

- Conference Automatic Software Engineering (ASE), ACM, pp. 525–526 (2007). <https://doi.org/10.1145/1321631.1321726>
- Li, Z., Zhou, Y.: PR-Miner: Automatically Extracting Implicit Programming Rules and Detecting Violations in Large Software Code. In: Proceedings 5th European Software Engineering Conference/Found Software Engineering (ESEC/FSE), ACM, pp. 306–315 (2005). <https://doi.org/10.1145/1095430.1081755>
- Livshits, B., Zimmermann, T.: DynaMine: Finding Common Error Patterns by Mining Software Revision Histories. In: Proceedings 5th European Software Engineering Conference /Found Software Engineering (ESEC/FSE), ACM, vol 30, pp. 296–305 (2005). <https://doi.org/10.1145/1095430.1081754>
- Long, F., Rinard, M.: Automatic Patch Generation by Learning Correct Code. In: Proceedings 43rd Symposium Principles Program Language (POPL), ACM, New York, NY, USA, pp. 298–312 (2016). <https://doi.org/10.1145/2837614.2837617>
- Lv, F., Zhang, H., Lou, J.g., Wang, S., Zhang, D., Zhao, J.: CodeHow: Effective Code Search Based on API Understanding and Extended Boolean Model. In: Proceedings 30th International Conference Automatic Software Engineering (ASE), IEEE Press, pp. 260–270 (2015). <https://doi.org/10.1109/ASE.2015.42>
- Mandelin, D., Xu, L., Bodík, R., Kimelman, D.: Jungloid Mining: Helping to Navigate the API Jungle. In: Proceedings 26th ACM SIGPLAN Conference Programming Language Design Implementation (PLDI), ACM, New York, NY, USA, pp. 48–61 (2005). <https://doi.org/10.1145/1065010.1065018>
- McMillan, C., Grechanik, M., Poshvanyk, D., Xie, Q., Fu, C.: Portfolio: Finding Relevant Functions and Their Usage. In: Proceedings 33rd International Conference Software Engineering (ICSE), ACM, New York, NY, USA, pp. 111–120 (2011). <https://doi.org/10.1145/1985793.1985809>
- Mockus, A., Weiss, D.M.: Predicting risk of software changes. *Bell Labs Tech. J.* 5(2), 169–180 (2000). <https://doi.org/10.1002/bltj.2229>
- Moreno, L., Bavota, G., Di Penta, M., Oliveto, R., Marcus, A.: How Can I Use This Method? In: Proceedings 37th International Conference Software Engineering (ICSE), IEEE, vol 1, pp. 880–890 (2015). <https://doi.org/10.1109/ICSE.2015.98>
- Murali, V., Chaudhuri, S., Jermaine, C.: Bayesian Specification Learning for Finding API Usage Errors. In: Proceedings 11th European Software Engineering Conference/Found Software Engineering (ESEC/FSE), ACM, pp. 151–162 (2017). <https://doi.org/10.1145/3106237.3106284>
- Nadi, S., Krüger, S., Mezini, M., Bodden, E.: Jumping through Hoops: Why Do Java Developers Struggle with Cryptography APIs? In: Proceedings of the 38th International Conference on Software Engineering (ICSE), ACM, New York, NY, USA, pp. 935–946 (2016). <https://doi.org/10.1145/2884781.2884790>
- Nguyen, T.D., Nguyen, A.T., Phan, H.D., Nguyen, T.N.: Exploring API Embedding for API Usages and Applications. In: Proceedings 39th International Conference Software Engineering (ICSE), IEEE, pp. 438–449 (2017). <https://doi.org/10.1109/ICSE.2017.47>
- Nguyen, T.T., Nguyen, H.A., Pham, N.H., Al-Kofahi, J.M., Nguyen, T.N.: Graph-based mining of multiple object usage patterns. In: Proceedings 7th European Software Engineering Conference/Found Software Engineering (ESEC/FSE), ACM, New York, NY, USA, pp. 383–392 (2009). <https://doi.org/10.1145/1595696.1595767>
- Nielebock, S., Heumüller, R., Ortmeier, F.: Commits As a Basis for API Misuse Detection. In: Proceedings 7th International Work on Software Min ACM, New York, NY, USA, pp. 20–23 (2018). <https://doi.org/10.1145/3242887.3242890>
- Nielebock, S., Heumüller, R., Krüger, J., Ortmeier, F.: Cooperative API Misuse Detection Using Correction Rules. In: ACM (ed.) Proceedings 42nd International Conference Software Engineering (ICSE), ACM (2020a). <https://doi.org/10.1145/3377816.3381735>, ICSE-NIER track - forthcoming
- Nielebock, S., Heumüller, R., Krüger, J., Ortmeier, F.: Using API-Embedding for API-Misuse Repair. In: ACM (ed.) Proceedings 1st International Work Automation Program Repair (APR) (2020b). <https://doi.org/10.1145/3387940.3392171>, forthcoming
- Nielebock, S., Blockhaus, P., Krüger, J., Ortmeier, F.: AndroidCompass: A Dataset of Android Compatibility Checks in Code Repositories. In: Proceedings 18th International Work Min Software Repository (MSR), pp. 535–539 (2021). <https://doi.org/10.1109/MSR52588.2021.00069>
- Oliveira, D.S., Lin, T., Rahman, M.S., Akefirad, R., Ellis, D., Perez, E., Bobhate, R., DeLong, L.A., Cappos, J., Brun, Y.: API Blindspots: Why Experienced Developers Write Vulnerable Code. In: Proceedings 14th Symposium on Usable Privacy and Security (SOUPS), USENIX Association, Baltimore, MD, pp. 315–328 (2018)

- Pasquier, N., Bastide, Y., Taouil, R., Lakhal, L.: Discovering Frequent Closed Itemsets for Association Rules. In: Beeri, C., Buneman, P. (eds.) *Proceedings 7th International Conference Database Theory (ICDT)*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 398–416 (1999). https://doi.org/10.1007/3-540-49257-7_25
- Paul, S., Prakash, A.: A framework for source code search using program patterns. *IEEE Trans. Softw. Eng.* **20**(6), 463–475 (1994). <https://doi.org/10.1109/32.295894>
- Pradel, M., Gross, T.R.: Automatic Generation of Object Usage Specifications from Large Method Traces. In: *Proceedings 24th International Conference Automation Software Engineering (ASE)*, IEEE Computer Society, pp. 371–382 (2009). <https://doi.org/10.1109/ASE.2009.60>
- Rahman, M.M., Roy, C.K., Lo, D.: RACK: Automatic API Recommendation Using Crowdsourced Knowledge. In: *Proceedings 23rd International Conference Software Analytical Evolution, Reeng. (SANER)*, IEEE, vol 1, pp. 349–359 (2016). <https://doi.org/10.1109/SANER.2016.80>
- Rahman, M.M., Barson, J., Paul, S., Kayani, J., Lois, F.A., Quezada, S.F., Parnin, C., Stolee, K.T., Ray, B.: Evaluating How Developers Use General-Purpose Web-Search for Code Retrieval. In: *Proceedings of the 15th International Conference on Mining Software Repositories (MSR)*, ACM, New York, NY, USA, pp. 465–475 (2018). <https://doi.org/10.1145/3196398.3196425>
- Reiss, S.P.: Semantics-Based Code Search. In: *Proceedings 31st International Conference Software Engineering (ICSE)*, IEEE Computer Society, pp. 243–253 (2009). <https://doi.org/10.1109/ICSE.2009.5070525>
- Robillard, M.P., Deline, R.: A field study of API learning obstacles. *Empir. Softw. Eng.* **16**(6), 703–732 (2011). <https://doi.org/10.1007/s10664-010-9150-8>
- Robillard, M.P., Bodden, E., Kawrykow, D., Mezini, M., Ratchford, T.: Automated API property inference techniques. *IEEE Trans. Softw. Eng.* **39**(5), 613–637 (2013). <https://doi.org/10.1109/TSE.2012.63>
- Roy, C.K., Cordy, J.R., Koschke, R.: Comparison and evaluation of code clone detection techniques and tools: a qualitative approach. *Sci. Comput. Progr.* **74**(7), 470–495 (2009). <https://doi.org/10.1016/j.scico.2009.02.007>
- Sadowski, C., Stolee, K.T., Elbaum, S.: How Developers Search for Code: A Case Study. In: *Proceedings 10th European Software Engineering Conference/Found Software Engineering (ESEC/FSE)*, ACM, pp. 191–201 (2015). <https://doi.org/10.1145/2786805.2786855>
- Sahavechaphan, N., Claypool, K.: XSnippet: Mining For Sample Code. In: *Proceedings 21st Conference Object-Oriented Program., System, Language, Application. (OOPSLA)*, ACM, New York, NY, USA, pp. 413–430 (2006). <https://doi.org/10.1145/1167473.1167508>
- Saied, M.A., Raelijohn, E., Batot, E., Famelis, M., Sahraoui, H.: Towards assisting developers in API usage by automated recovery of complex temporal patterns **119**, 106213 (2020). <https://doi.org/10.1016/j.infsof.2019.106213>
- Saini V, Farmahinifarahani F, Lu Y, Baldi P, Lopes CV.: Oreo: Detection of Clones in the Twilight Zone. In: *Proceedings 26th European Software Engineering Conference/Found Software Engineering (ESEC/FSE)*, ACM, New York, NY, USA, pp. 354–365 (2018). <https://doi.org/10.1145/3236024.3236026>
- Sajani, H., Saini, V., Svajlenko, J., Roy, C.K., Lopes, C.V.: SourcererCC: Scaling Code Clone Detection to Big-Code. In: *Proceedings 38th International Conference Software Engineering (ICSE)*, IEEE, pp. 1157–1168 (2016). <https://doi.org/10.1145/2884781.2884877>
- Saul, Z.M., Filkov, V., Devanbu, P., Bird, C. Recommending Random Walks. In: *Proceedings 6th European Software Engineering Conference/Found. Software Engineering (ESEC/FSE)*, ACM, New York, NY, USA, pp. 15–24 (2007). <https://doi.org/10.1145/1287624.1287629>
- Siegmund, J., Siegmund, N., Apel, S.: Views on Internal and External Validity in Empirical Software Engineering. In: *Proceedings 37th International Conference Software Engineering (ICSE)*, vol 1, pp. 9–19 (2015). <https://doi.org/10.1109/ICSE.2015.24>
- Śliwinski, J., Zimmermann, T., Zeller, A.: When Do Changes Induce Fixes? In: *Proceedings 2nd International Work Min Software Repository (MSR)*, ACM, New York, NY, USA, pp. 1–5 (2005). <https://doi.org/10.1145/1082983.1083147>
- Stolee, K.T., Elbaum, S., Dobos, D.: Solving the search for source code. *ACM Trans. Softw. Eng. Methodol.* **23**(3), (2014). <https://doi.org/10.1145/2581377>
- Sun, B., Shu, G., Podgurski, A., Li, S., Zhang, S., Yang, J.: Propagating Bug Fixes with Fast Subgraph Matching. In: *Proceedings 21st International Symposium Software Reliable Engineering (ISSRE)*, pp. 21–30 (2010). <https://doi.org/10.1109/ISSRE.2010.36>
- Thummalapenta, S., Xie, T.: Parseweb: A Programmer Assistant for Reusing Open Source Code on the Web. In: *Proceedings 22nd International Conference Automation Software Engineering (ASE)*, ACM, New York, NY, USA, pp. 204–213 (2007). <https://doi.org/10.1145/1321631.1321663>

- Thung, F., Lo, D., Lawall, J.: Automated Library Recommendation. In: Proceedings 20th Work Conference Reverse Engineering (WCRE), pp. 182–191 (2013a). <https://doi.org/10.1109/WCRE.2013.6671293>
- Thung, F., Wang, S., Lo, D., Lawall, J.: Automatic Recommendation of API Methods from Feature Requests. In: Proceedings 28th International Conference Automation Software Engineering (ASE), IEEE Press, pp. 290–300 (2013b). <https://doi.org/10.1109/ASE.2013.6693088>
- Wasylkowski, A., Zeller, A.: Mining temporal specifications from object usage. *Autom. Softw. Eng.* **18**(3–4), 263–292 (2011). <https://doi.org/10.1007/s10515-011-0084-1>
- Wasylkowski, A., Zeller, A., Lindig, C.: Detecting Object Usage Anomalies. In: Proceedings 6th European Software Engineering Conference/Found. Software Engineering (ESEC/FSE), ACM, pp. 35–44 (2007). <https://doi.org/10.1145/1287624.1287632>
- Weimer, W., Necula, G.C.: Mining Temporal Specifications for Error Detection. In: Halbwachs, N., Zuck, L.D. (eds.) Proceedings 11th International Conference Tools Algorithms Construct Analytical System (TACAS), Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 461–476 (2005). https://doi.org/10.1007/978-3-540-31980-1_30
- White, M., Tufano, M., Vendome, C., Poshvanyk, D.: Deep Learning Code Fragments for Code Clone Detection. In: Proceedings 31st International Conference Automation Software Engineering (ASE), ACM, pp. 87–98 (2016). <https://doi.org/10.1145/2970276.2970326>
- Wu, R., Wen, M., Cheung, S.C., Zhang, H.: ChangeLocator: locate crash-inducing changes based on crash reports. *Empir. Softw. Eng.* (2017). <https://doi.org/10.1007/s10664-017-9567-4>
- Xia, X., Bao, L., Lo, D., Kochhar, P.S., Hassan, A.E., Xing, Z.: What do developers search for on the web? *Empir. Softw. Eng.* **22**(6), 3149–3185 (2017). <https://doi.org/10.1007/s10664-017-9514-4>
- Xin, Q., Reiss, S.P.: Leveraging Syntax-Related Code for Automated Program Repair. In: Proceedings 32nd International Conference Automation Software Engineering (ASE), IEEE Press, pp. 660–670 (2017). <https://doi.org/10.1109/ASE.2017.8115676>
- Xin, Q., Reiss, S.P.: Better Code Search and Reuse for Better Program Repair. In: Proceedings 6th, IEEE Press, pp. 10–17 (2019). <https://doi.org/10.1109/GI.2019.00012>
- Yang, J., Evans, D.: Automatically Inferring Temporal Properties for Program Evolution. In: Proceedings 15th International Symposium Software Reliable Engineering (ISSRE), IEEE, pp. 340–351 (2004). <https://doi.org/10.1109/ISSRE.2004.11>
- Yang, J., Evans, D., Bhardwaj, D., Bhat, T., Das, M.: Perracotta: Mining Temporal API Rules from Imperfect Traces. In: Proceedings 28th International Conference Software Engineering (ICSE), ACM, pp. 282–291 (2006). <https://doi.org/10.1145/1134285.1134325>
- Yates, F.: Contingency tables involving small numbers and the chi2 test. *Suppl. J. R. Stat. Soc.* **1**(2), 217–235 (1934). <https://doi.org/10.2307/2983604>
- Zhong, H., Su, Z.: An Empirical Study on Real Bug Fixes. In: Proceedings 37th International Conference Software Engineering (ICSE), IEEE/ACM, IEEE Press, pp. 913–923 (2015). <https://doi.org/10.1109/ICSE.2015.101>
- Zhong, H., Xie, T., Zhang, L., Pei, J., Mei, H.: MAPO: Mining and Recommending API Usage Patterns. Proceedings 23rd European Conference Object-Oriented Program (ECOOP), pp. 318–343 (2009). https://doi.org/10.1007/978-3-642-03013-0_15
- Zhou, Y., Gu, R., Chen, T., Huang, Z., Panichella, S., Gall, H.: Analyzing APIs Documentation and Code to Detect Directive Defects. In: Proceedings 39th International Conference Software Engineering (ICSE), pp. 27–37 (2017). <https://doi.org/10.1109/ICSE.2017.11>
- Zibran, M.F., Eishita, F.Z., Roy, C.K.: Useful, But Usable? Factors Affecting the Usability of APIs. In: Proceedings 18th Work Conference Reverse Engineering (WCRE), IEEE, pp. 151–155 (2011). <https://doi.org/10.1109/WCRE.2011.26>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Sebastian Nielebock¹ · Robert Heumüller¹ · Kevin Michael Schott¹ · Frank Ortmeier¹

Robert Heumüller
robert.heumueller@ovgu.de

Kevin Michael Schott
kevin.schott@ovgu.de

Frank Ortmeier
frank.ortmeier@ovgu.de

- ¹ Faculty of Computer Science, Chair of Software Engineering, Otto-von-Guericke University Magdeburg, Magdeburg, Germany