

Augmenting Stack Overflow with API Usage Patterns Mined from GitHub

Anastasia Reinhardt^{1,2*} Tianyi Zhang^{1†} Mihir Mathur¹ Miryung Kim¹

¹University of California, Los Angeles

²George Fox University

areinhardt14@georgefox.edu, {tianyi.zhang, miryung}@cs.ucla.edu, mihirmathur@ucla.edu

ABSTRACT

Programmers often consult Q&A websites such as Stack Overflow (SO) to learn new APIs. However, online code snippets are not always complete or reliable in terms of API usage. To help programmers assess online code snippets, we build a Chrome-extension, EXAMPLECHECK that detects API usage violations in SO posts using API usage patterns mined from 380K GitHub projects. It quantifies how many GitHub examples follow common API usage and illustrates how to remedy the detected violation in a given SO snippet. With EXAMPLECHECK, programmers can easily identify the pitfalls of a given SO snippet and learn how much it deviates from common API usage patterns found in GitHub. EXAMPLECHECK is available at Chrome Web Store [2] and the demo video is available at <https://youtu.be/WOnN-wQZsH0>.

CCS CONCEPTS

• **Software and its engineering** → *Software reliability; Collaboration in software development; Integrated and visual development environments;*

KEYWORDS

online Q&A forum, API usage pattern, code assessment

ACM Reference Format:

Anastasia Reinhardt, Tianyi Zhang, Mihir Mathur, and Miryung Kim. 2018. Augmenting Stack Overflow with API Usage Patterns Mined from GitHub. In *Proceedings of the 26th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '18)*, November 4–9, 2018, Lake Buena Vista, FL, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3236024.3264585>

1 INTRODUCTION

Programmers often search for online code examples to learn new APIs. A case study at Google shows that developers issue an average of 12 code search queries per weekday [8]. Stack Overflow (SO) is a popular Q&A website that programmers often resort to. As of July 2017, Stack Overflow has accumulated more than 22

*Work done as an intern at University of California, Los Angeles.

†Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ESEC/FSE '18, November 4–9, 2018, Lake Buena Vista, FL, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5573-5/18/11...\$15.00

<https://doi.org/10.1145/3236024.3264585>

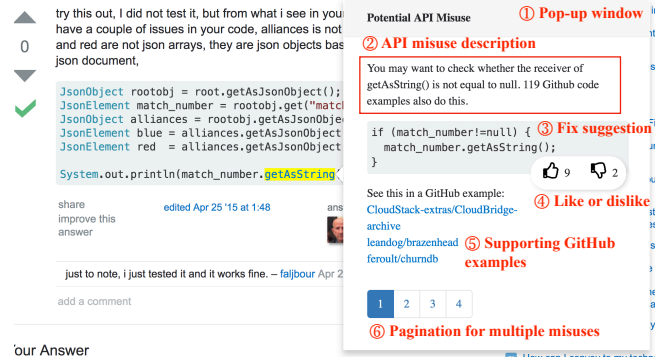


Figure 1: The EXAMPLECHECK Chrome extension that augments Stack Overflow with API misuse warnings. The pop-up window alerts that `match_number` can be null if the requested JSON attribute does not exist and will crash the program by throwing `NullPointerException` when `getAsString` is called on it.

million answers, many of which contain code snippets for specific programming questions. However, SO snippets are not always complete or reliable, which can be misleading and sometimes harmful when programmers follow them *as-is* during software development. For example, Fischer et al. find that 29% of security-related code snippets in Stack Overflow are insecure and might affect over 1 million Android apps in Google play [7].

This tool demonstration paper builds on the API usage mining and API misuse detection technique described in our ICSE 2018 paper [11]. Our insight is that common API usage inferred from a large corpus of 380K GitHub projects may represent a desirable pattern that a programmer can use to examine and enhance SO code snippets. Mined API usage patterns abstract away syntactic details such as variable names, but retain the temporal ordering, control structures, and guard conditions of API calls.

This paper, in particular, focuses on the tool features and implementation details of a Chrome extension, called EXAMPLECHECK that informs programmers about API usage violations in SO posts. Figure 1 shows a screenshot of EXAMPLECHECK. Given a SO post, EXAMPLECHECK first extracts the sequence of API calls with corresponding control constructs and guard conditions. EXAMPLECHECK then contrasts the call sequence against the common API usage patterns mined from GitHub. To help users understand a detected violation, EXAMPLECHECK generates a descriptive warning message and illustrates how to fix the violation with corresponding GitHub examples. Using common API usage patterns to detect violations may lead to false alarms, since these mined, common patterns do

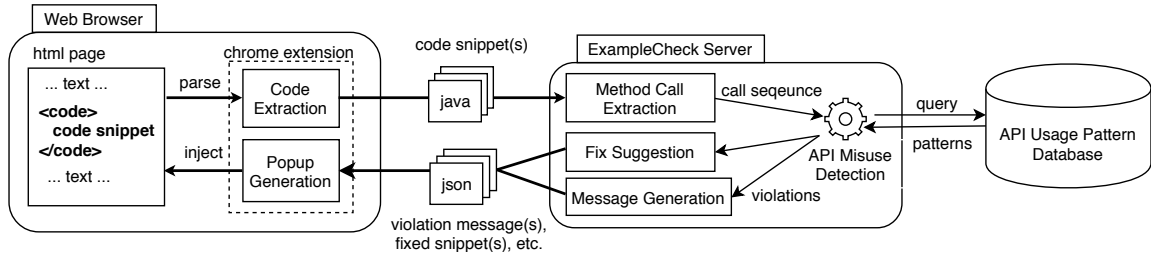


Figure 2: An overview of EXAMPLECHECK's architecture

not necessarily represent correct API usage. To mitigate this issue, EXAMPLECHECK allows users to upvote or downvote a reported violation based on its applicability and usefulness.

The resulting data set of mined API usage patterns, detected violations in Stack Overflow, and our manual inspection result are publicly available at <http://web.cs.ucla.edu/~tianyi.zhang/examplecheck.html>. The Chrome extension is available for download at Chrome Web Store [2].

2 APPROACH & IMPLEMENTATION

This section describes the tool implementation details of EXAMPLECHECK. Figure 2 shows the architecture of EXAMPLECHECK. The API usage mining process is computed offline and the resulting patterns are stored in a database. The technical details and evaluation of API usage mining technique is presented in our ICSE 2018 paper [11]. When a user loads a Stack Overflow page in the Chrome browser, the Chrome extension extracts code snippets within `<code>` tags in answer posts, and sends them to the back-end server. The back end then detects API usage violations in a snippet and synthesizes warning messages and corresponding fixes. For each misused method call in a snippet, the Chrome extension generates a pop-up window using the Bootstrap popover plug-in¹ to inform the user about the API misuse information.

API Usage Mining and the Resulting Pattern Set. Our mining technique in [11] leverages a distributed software mining infrastructure [6] to search over the corpus of 380K GitHub projects. Given an API method of interest, it identifies code fragments that use the same method in the GitHub corpus and performs program slicing to remove statements that are not related to the given method. Then it combines frequent subsequence mining and SMT-based guard condition mining to retain important API usage features, including the temporal ordering of related API calls, enclosing control structures, and guard conditions that protect an API call. We evaluated the mining technique using 30 API methods from MUBench [3]. Our mining technique has 80% precision and 91% recall, when considering top 5 patterns for each API method.

In our prior work [11], we mined API usage patterns of 100 popular Java API methods and carefully inspected 245 inferred patterns based on online documentation. As a result, we curated a dataset of 180 validated, correct patterns for API misuse detection, which covers API usages shown in 217K SO posts in Java. These patterns are represented as API call sequences with surrounding control constructs. Each API call is also annotated with its argument

types and guard conditions. For example, one pattern, `loop { ; get(int)@arg0<rcv.size(); } ;`, checks if the index is out of bounds when calling the `get` method on an `ArrayList` object.

API Misuse Detection. Given a code snippet sent from the browser, the server first extracts the API call sequence from the snippet. We use a partial program analysis and type resolution technique [9] to parse incomplete snippets and resolve ambiguous types. If a SO snippet has multiple methods, EXAMPLECHECK inlines the call sequence of an invoked method into the sequence of the caller to emulate a lightweight inter-procedural analysis. EXAMPLECHECK then queries the pattern database for the API calls present in each API call sequence. Given an API call sequence and an API usage pattern, it checks whether (1) the API calls and control constructs in the sequence follow the same temporal order in the pattern, and (2) the guard condition of an API call in the sequence implies the guard of the corresponding API call in the pattern. EXAMPLECHECK uses a SMT solver, Z3 [5], to check whether one guard condition implies another. EXAMPLECHECK is capable of detecting three types of API usage violations—*missing control constructs*, *missing or incorrect order of API call*, and *incorrect guard condition*.

Warning Message Generation. Given an API usage violation and the correct pattern, EXAMPLECHECK generates a warning message that describes the violation in natural language text. Table 1 shows the warning message templates for different types of API usage violations. In each template, `<?>` is instantiated with the corresponding API calls or control constructs based on the detected API usage violation and the correct pattern. `<before/after>` is instantiated based on the relative order of the two API calls in the correct pattern. The warning messages also describe which exception types are not handled in the snippets detected with *missing try-catch* violations. To help users understand the prevalence of a recommended API usage pattern, the warning message also quantifies how many other code fragments follow the same pattern in GitHub.

Fix Suggestion. EXAMPLECHECK further suggests a correct way of using an API method by synthesizing a readable fixed snippet based on the original SO snippet. EXAMPLECHECK first matches each API call in the recommended API usage pattern with the given SO snippet. If an API call is matched, EXAMPLECHECK reuses the same receiver object and arguments of the corresponding API call from the original SO snippet in the synthesized snippet. Otherwise, EXAMPLECHECK names the receiver and arguments based on their types. For example, if the receiver type of an unmatched API call (i.e., a *missing-API-call* violation) is `File`, EXAMPLECHECK names the receiver object as `file`, the lower case of the receiver type. In this

¹https://www.w3schools.com/bootstrap/bootstrap_popover.asp

Violation Type	Description Template	Example Warning Message
Missing/Incorrect Order of API calls	You may want to call <?> <before/after> calling <?>	You may want to call <code>TypedArray.recycle()</code> after calling <code>TypedArray.getString()</code> . [35784171]
Missing Control Constructs	You may want to call the API method <?> in <?>	You may want to call <code>Cursor.close()</code> in a finally block. [31427468]
Missing Try-Catch	You may want to handle the potential <?> exception thrown by <?> using a try-catch block	You may want to handle the potential <code>SQLException</code> thrown by <code>PreparedStatement.setString()</code> using a try-catch block. [11183042]
Incorrect Guard Conditions	You may want to check whether <?> is true before calling <?>	You may want to check whether <code>iterator.hasNext()</code> is true. [25789601]

Table 1: Warning message templates for different types of API usage violations. <?> and <before/after> are instantiated based on API usage violations and correct patterns. The digits in the last column are the SO post ids of the warning examples.

way, EXAMPLECHECK reduces the mental gap for switching between the original SO post and the recommended snippet.

3 DEMONSTRATION SCENARIO

Suppose Alice wants to read attribute values from a JSON message using Google’s Gson library. Alice searches online and finds a related Stack Overflow post with an illustrative code example, as shown in Figure 1.² Though this post is accepted as a correct answer, it does not properly use the `JsonElement.getAsString` method, which gets the string value of a JSON element. For example, if the requested attribute does not exist in the JSON message, the preceding API call, `JsonObject.get` will return null, which consequently leads to `NullPointerException` when calling `getAsString` on the returned object. If Alice puts too much trust on this example of the SO post, she may inadvertently follow an unreliable solution, which might lead to runtime errors in some corner cases.

Alice cannot easily recognize the potential limitation of the given SO post, unless she manually investigates other similar code examples. EXAMPLECHECK frees Alice from this manual investigation labor by contrasting a Stack Overflow post with common API usage patterns mined from over 380K GitHub repositories. EXAMPLECHECK then highlights the potential API usage violations in the Stack Overflow post. When Alice clicks on a highlighted API call, EXAMPLECHECK generates a pop-up window with detailed descriptions about the API usage violation, as shown in Figure 1.

API misuse description. To help Alice understand a detected API usage violation, EXAMPLECHECK translates the violation to a natural language description (① in Figure 1). From the warning message, Alice learns that she should check whether the `JsonElement` object is null before calling `getAsString`. Furthermore, EXAMPLECHECK displays a message that 119 GitHub examples also follow this API usage pattern. Such quantification can provide additional evidence about how many real-world examples are different from the given SO snippet.

Fix suggestion. EXAMPLECHECK further sketches how to correct the violation in the original SO post, as shown in ③ in Figure 1. This fix is an embodiment of the correct API usage pattern in the context of the SO post. To reduce the gap between the fix and the original post, EXAMPLECHECK reuses the same variable names in the original SO posts to generate a suggestion with improved API usage. For example, the `JsonElement` variable in the generated example is named as the same variable, `match_number` in the original post.

Linking GitHub examples. To help Alice understand how the same API method is used in real-world projects, EXAMPLECHECK



Figure 3: EXAMPLECHECK will redirect a programmer to a concrete code example from GitHub that follows a correct API usage pattern, when she clicks on a GitHub example link in the pop-up window.³

provides several GitHub examples that follow the suggested API usage pattern (⑤ in Figure 1). Alice is curious about how others use `JsonElement.getAsString`. When she clicks on the link of the first GitHub example, EXAMPLECHECK redirects Alice to a GitHub page and automatically scrolls down to the Java method where `JsonElement.getAsString` is called, as shown in Figure 3. Compared with the simplified SO example in Figure 1, this GitHub code is more carefully constructed with multiple if checks. For example, it not only checks whether the `JsonElement` object is null, but also checks whether it is a primitive type to avoid `ClassCastException` before calling `getAsString`. By providing the traceability to concrete code examples in GitHub, Alice could gain a more comprehensive view of correct API usage in production code, which may not be illustrated in simplified code examples in Stack Overflow.

User feedback. After investigating the concrete example in GitHub, Alice feels that it is necessary to perform a null check. She upvotes the pattern by clicking on the “thumbs-up” button to notify other users that this detected violation is helpful to her (④ in Figure 1). Alice also finds that her decision resonates with the majority of EXAMPLECHECK users, since nine other users also upvoted this detected violation.

²<https://stackoverflow.com/questions/29860000>

³<https://goo.gl/YHo1UM>

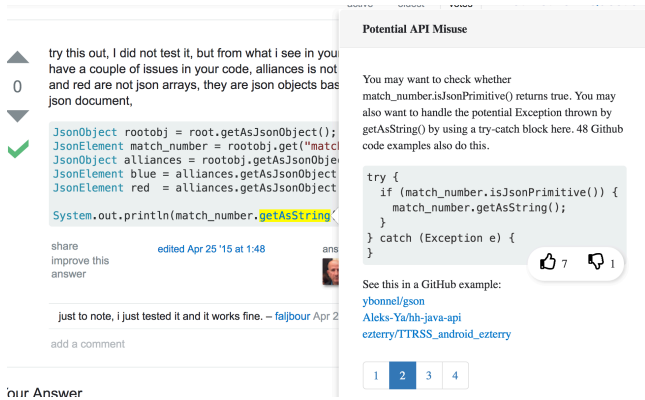


Figure 4: Another API usage warning that reminds programmers to check whether the `JsonElement` object represents a JSON primitive value by calling `isJsonPrimitive`. It also suggests to catch potential exceptions thrown by `getString`.

Multiple API usage violations. If a method call in a SO post violates multiple API usage patterns, EXAMPLECHECK displays them in separate pages in a pop-up window. These pages are first ranked by the vote score of each violated pattern, e.g., upvotes minus downvotes. If two patterns have the same vote score, EXAMPLECHECK then ranks them by the number of GitHub examples that support a pattern. As shown in ⑥ in Figure 1, the method call, `getString` violates four API usage patterns. Figure 4 shows the second violated pattern and suggests Alice to check whether the `JsonElement` object represents a JSON primitive value before calling the `getString` method. Otherwise, `getString` will throw `ClassCastException`. EXAMPLECHECK also suggests Alice to wrap `getString` with a try-catch block to handle potential exceptions. This pattern is supported by 48 GitHub examples.

4 RELATED WORK

Prior work has investigated the quality of code snippets in Stack Overflow. Several studies show that SO snippets are often incomplete and the API names appearing in these snippets are hard to resolve [4, 9]. Zhou et al. observe that 86 of 200 accepted SO posts use deprecated APIs but only 3 of them are reported by other users [12]. Fischer et al. find that 29% of security-related SO snippets are insecure and have potentially been reused to over 1 million Android apps on Google play [7]. Treude and Robillard conduct a survey to investigate comprehension difficulty of code examples in Stack Overflow [10]. The responses from GitHub users indicate that less than half of the SO examples are self-explanatory due to issues such as incomplete code and missing explanations. Though we draw motivation from these studies, EXAMPLECHECK focuses on detecting API usage violations by contrasting SO code examples against common API usage patterns mined from GitHub. While EXAMPLECHECK follows a similar style to Codota [1], Codota does not group related examples based on common API usage, does not quantify how many GitHub code snippets support the common usage, and does not detect API misuse by contrasting the SO post against desirable API usage mined from GitHub.

5 SUMMARY

The main contribution of this paper is the design and implementation of EXAMPLECHECK, which provides browser-based tool support for systematically assessing and augmenting Stack Overflow with common API usage patterns mined from GitHub. In our previous work [11], we examine 217K SO posts with 180 validated patterns and find that 31% of SO posts have potential API usage violations. The first two authors manually inspect 400 SO posts with detected API usage violations and confirm real API misuse in 289 posts, which can produce symptoms such as program crashes and resource leaks if the posts are reused *as-is* to target projects. We also find that many unreliable examples are simplified to operate on crafted input data for illustration purposes only. Such curated examples could be insufficient for various input data and usage scenarios in real software systems, especially for handling corner cases.

As future work, we plan to conduct a longitudinal study with Stack Overflow users to understand the adoption and usage of EXAMPLECHECK. We will also investigate different weighting schemes for effectively ranking detected API usage violations.

6 ACKNOWLEDGMENT

Participants in this project are in part supported through AFRL grant FA8750-15-2-0075, NSF grants CCF-1764077, CCF-1527923, CCF-1460325, CCF-1723773, ONR grant N00014-18-1-2037 and gift from Huawei.

REFERENCES

- [1] 2018. Codota Code Browsing Assistant. <https://www.codota.com/code-browsing-assistant>.
- [2] 2018. ExampleCheck - Chrome Web Store. <https://chrome.google.com/webstore/detail/examplecheck/amliempebckaiaiklimcpomlnkikioe>.
- [3] Sven Amani, Sarah Nadi, Hoan A Nguyen, Tien N Nguyen, and Mira Mezini. 2016. MUBench: a benchmark for API-misuse detectors. In *Proceedings of the 13th International Conference on Mining Software Repositories*. ACM, 464–467.
- [4] Barthélemy Dagenais and Martin P Robillard. 2012. Recovering traceability links between an API and its learning resources. In *Proceedings of the 34th International Conference on Software Engineering*. IEEE, 47–57.
- [5] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.
- [6] Robert Dyer, Hoan Anh Nguyen, Hridesh Rajan, and Tien N Nguyen. 2013. Boa: A language and infrastructure for analyzing ultra-large-scale software repositories. In *Proceedings of the 35th International Conference on Software Engineering*. IEEE, 422–431.
- [7] Felix Fischer, Konstantin Böttinger, Huang Xiao, Christian Stransky, Yasemin Acar, Michael Backes, and Sascha Fahl. 2017. Stack overflow considered harmful? the impact of copy&paste on android application security. In *2017 IEEE Symposium on Security and Privacy*. IEEE, 121–136.
- [8] Caitlin Sadowski, Kathryn T Stolee, and Sebastian Elbaum. 2015. How developers search for code: a case study. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*. ACM, 191–201.
- [9] Siddharth Subramanian, Laura Inozemtseva, and Reid Holmes. 2014. Live API documentation. In *Proceedings of the 36th International Conference on Software Engineering*. ACM, 643–652.
- [10] Christoph Treude and Martin P Robillard. 2017. Understanding Stack Overflow Code Fragments. In *Proceedings of the 33rd International Conference on Software Maintenance and Evolution*. IEEE, 509–513.
- [11] Tianyi Zhang, Ganesha Upadhyaya, Anastasia Reinhardt, Hridesh Rajan, and Miryung Kim. 2018. Are code examples on an online Q&A forum reliable?: a study of API misuse on stack overflow. In *Proceedings of the 40th International Conference on Software Engineering*. ACM, 886–896.
- [12] Jing Zhou and Robert J Walker. 2016. API deprecation: a retrospective analysis and detection method for code examples on the web. In *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 266–277.