

What Do Contribution Guidelines Say About Software Testing?

Bruna Falcucci, Felipe Gomide, Andre Hora
Department of Computer Science, UFMG
Belo Horizonte, Brazil
{brunafalcucci, felipe.gomide, andrehora}@dcc.ufmg.br

Abstract—Software testing plays a crucial role in the contribution process of open-source projects. For example, contributions introducing new features are expected to include tests, and contributions with tests are more likely to be accepted. Although most real-world projects require contributors to write tests, the specific testing practices communicated to contributors remain unclear. In this paper, we present an empirical study to understand better how software testing is approached in contribution guidelines. We analyze the guidelines of 200 Python and JavaScript open-source software projects. We find that 78% of the projects include some form of test documentation for contributors. Test documentation is located in multiple sources, including CONTRIBUTING files (58%), external documentation (24%), and README files (8%). Furthermore, test documentation commonly explains how to run tests (83.5%), but less often provides guidance on how to write tests (37%). It frequently covers unit tests (71%), but rarely addresses integration (20.5%) and end-to-end tests (15.5%). Other key testing aspects are also less frequently discussed: test coverage (25.5%) and mocking (9.5%). We conclude by discussing implications and future research.

Index Terms—Software Testing, Contribution Guidelines, Empirical Studies, GitHub

I. INTRODUCTION

Open-source software projects typically provide *contribution guidelines* detailing how to submit contributions [1]–[3]. For maintainers, contribution guidelines communicate how developers should contribute. For contributors, these guidelines support the verification of well-formed and valuable contributions. Both project maintainers and contributors benefit from saving time and avoiding problems caused by inadequately submitted contributions [1]. This kind of documentation may also alleviate barriers faced by new contributors [4]. Due to its importance, GitHub recommends having contribution guidelines as a prerequisite for launching open source projects [3] and provides guides to support their creation [1], [2].

Contribution guidelines may have multiple contents detailing the workflow, acceptance criteria, or how to create forks, branches, pull requests, and more [2]. Due to its importance to software development, software testing is also a fundamental part of the contributing process [2], [5]–[7]. Indeed, 52% of the contribution guidelines state that contributions have to include test cases [2]. Contributions introducing new features are expected to include tests [6], and core maintainers look for the inclusion of test cases as a signal of the thoroughness of the contribution [7]. Consequently, contributions that include test cases are more likely to be accepted [5].

GitHub recommends that contribution guidelines might require tests to improve the quality: “*Tests help contributors feel confident that they won’t break anything. They also make it easier for you to review and accept contributions quickly*” [3]. Similarly, popular guides with best practices recommend that contributions should at least detail how to run the tests [8]. It is not surprising that open-source projects treat tests seriously in their contribution guidelines. For example, the Pandas contribution documentation states “*Ensure you have appropriate tests. These should be the first part of any PR*”,¹ while the Flask one notes: “*Include tests if your patch adds or changes code*”.² Other projects go a step further by providing detailed information about testing in contribution guidelines. For instance, Cypress provides an overview of the test types contributors may create, such as unit, integration, and e2e.³ The CPython project has extensive contribution documentation on how to run/write the tests and how to increase test coverage.⁴

Despite software testing’s critical role in contribution guidelines, we are not yet aware of how software testing is explained to contributors. While most real-world projects require developers to write tests in contributions [2], it is unclear whether these projects actually communicate testing practices to support contributions. Furthermore, in the case projects do communicate, it is unclear what testing practices are presented for contributors. For example, maintainers may focus on various topics, including run/write tests, test types, coverage, mocking, and other software testing practices. Assessing this information can provide the basis for unveiling communication patterns (that should be adopted in contribution guidelines) and anti-patterns (that should be avoided).

In this paper, we present an empirical study to understand better how software testing is approached in the contribution guidelines of open-source software projects. We analyze the contribution guidelines of 200 popular open-source projects written in Python and JavaScript. We propose to answer the following research questions:

- **RQ1: How frequently is test documentation for contributors?** We find that 78% of the projects include

¹<https://pandas.pydata.org/docs/development/contributing.html>

²<https://github.com/pallets/flask/blob/main/CONTRIBUTING.rst>

³<https://github.com/cypress-io/cypress/blob/develop/CONTRIBUTING.md>

⁴<https://devguide.python.org/testing/run-write-tests>

some form of test documentation for contributors. Test documentation is mostly located in CONTRIBUTING files (58%) and external documentation (24%), but also occurs in README files (8%).

- **RQ2: What is the content of test documentation for contributors?** Test documentation commonly explains how to run tests (83.5%), but less often provides guidance on how to write tests (37%). Moreover, it frequently covers unit tests (71%), but rarely addresses integration (20.5%) and end-to-end tests (15.5%). Other key testing aspects are also less frequently discussed: test coverage (25.5%), mocking (9.5%), and best practices (9%).

Contributions. The contributions of this paper are twofold. First, we provide an initial empirical study to assess test documentation for contributors. Second, we discuss practical implications for researchers and practitioners.

II. STUDY DESIGN

A. Case Study

We aim to study the contribution guidelines of real-world and relevant software projects. For a better perspective of the contribution guidelines landscape, we focus on projects written in the two most popular programming languages nowadays: Python and JavaScript. For each programming language, we relied on the GitHub Search tool (GHS) [9] to collect the top-100 most popular software projects hosted on GitHub according to the number of stars (a metric primarily adopted in the software mining literature as a proxy of popularity [10], [11]). In this process, we took special care to filter out non-software projects, such as tutorials, examples, and code samples. Our final dataset is composed of 200 software projects (100 in Python and 100 in JavaScript). On the median, the selected projects have 35.7K stars.

B. Detection of Contribution Guidelines

GitHub recommends to write contribution guidelines in CONTRIBUTING files [1], as in jQuery.⁵ Contribution guidelines may also be written directly in README files [12], as in Webpack.⁶ More comprehensive guidelines may be located in external documentation, as in Black⁷ and CPython.⁸

In this study, we consider all kinds of contribution guidelines as long as they are dedicated to contributors. We look for contribution guidelines (or links) in README and CONTRIBUTING files as well as other variations, such as DEVELOPMENT files. We verify not only the root directory of the projects but also the source code and test folders. Following these steps, we found that most projects have contribution guidelines (184 out of 200). This is expected, considering the relevance and popularity of the selected projects.

All selected projects, contribution guidelines, and extracted information are publicly available in our dataset at <https://doi.org/10.5281/zenodo.14046800>.

⁵<https://github.com/jquery/jquery/blob/main/CONTRIBUTING.md>

⁶<https://github.com/webpack/webpack/blob/main/test/README.md>

⁷<https://black.readthedocs.io/en/latest/contributing>

⁸<https://devguide.python.org>

C. Detection of Test Documentation for Contributors

Next, we searched the contribution guidelines for any test documentation. We define *test documentation for contributors* as any information related to software testing that is aimed at contributors and is included in contribution guidelines. Particularly, we focused on essential testing topics such as how to run or write tests, best practices, test coverage, mocking, and test types (*i.e.*, unit, integration, and end-to-end) [3], [8]. If no test documentation is found for a given project, a second author conducts a follow-up search to double-check and minimize the risk of false negatives. Following these steps, we detected that 156 out of 200 projects include test documentation for contributors.

D. Research Questions

We propose two RQs to assess the test documentation for contributors. First, we explore the *frequency* of test documentation for contributors. The rationale is to better understand the extent of the phenomenon under study and to determine whether any differences exist between Python and JavaScript. Second, we explore the *content* of test documentation for contributors. The rationale is to uncover which testing topics are addressed in contribution guidelines and determine which are the most common and lacking. Such information may help discover best practices and guidelines for creating test documentation for contributors.

III. RESULTS

A. RQ1: Frequency of test documentation for contributors

We find test documentation for contributors in 78% of the analyzed projects (156 out of 200), as detailed in Table I. Notice that Python has slightly more systems with test documentation (82 in 100) than JavaScript (74 in 100).

TABLE I
SUMMARY OF TEST DOCUMENTATION FOR CONTRIBUTORS.

	Python	JavaScript	Total
Analyzed Projects	100	100	200
• With Test Documentation (#)	82	74	156
• With Test Documentation (%)	82%	74%	78%

Table II presents the location of the 156 test documentation. We see that most test documentation for contributors is naturally located in the recommended CONTRIBUTING files (58%). Test documentation is also present in external documentation (24%) and README files (8%). We also find test documentation in other files (10%), such as *developers*, *development*, *testing*, and *workflow*. Overall, there is no major difference between Python and JavaScript regarding the location. The sole discrepancy happens in the location external documentation: Python has 26 in this location, while JavaScript has only 12. This may be explained by the fact that Python has the *Read the Docs*⁹ platform to build documentation (6 out of the 26 Python systems rely on it).

⁹<https://about.readthedocs.com>

TABLE II
LOCATION OF TEST DOCUMENTATION FOR CONTRIBUTORS.

Location	Python	JavaScript	Total	%
CONTRIBUTING file	41	49	90	58%
External documentation	26	12	38	24%
README file	5	7	12	8%
Other files	10	6	16	10%
All	82	74	156	100%

Finding 1: (i) 78% of the projects include some form of test documentation for contributors. (ii) Test documentation is mostly located in CONTRIBUTING files (58%) and external documentation (24%), but also occurs in README files (8%).

B. RQ2: Content of test documentation for contributors

Table III summarizes the content of the test documentation for contributors. It is divided into five categories: (1) how to run/write tests, (2) test types, (3) test coverage, (4) mocking, and (5) best practices & tips. The most frequent content is *how to run/write tests* (present in 86% of the test documentation), followed by *test types* (74%) and *test coverage* (25.5%). Next, we briefly describe each category.

TABLE III
CONTENT OF TEST DOCUMENTATION FOR CONTRIBUTORS.

Content	Python	JavaScript	Total	%
How to run/write tests	72	62	134	86%
• How to run tests	70	60	130	83.5%
• How to write tests	30	28	58	37%
Test types	65	50	115	74%
• Unit test	64	47	111	71%
• Integration test	20	12	32	20.5%
• e2e test	7	17	24	15.5%
Test coverage	31	9	40	25.5%
Mocking	10	5	15	9.5%
Best practices & tips	8	6	14	9%

How to run/write tests: Most real-world projects require contributors to test their patches locally and write new tests if code is added or changed [2], [3], [6]. Therefore, ideally, test documentation should describe at least how to run and write tests. Overall, we find that 86% (134 in 156) of the test documentation includes information about running and/or writing tests. While 83.5% describe how to run tests, only 37% actually detail how to write tests. The typical how-to-run test simply includes running commands, as in the test documentation of Meteor.¹⁰ Some test documentations go a step further, explaining how to write tests as in Next.js¹¹ and

¹⁰<https://github.com/meteor/meteor/blob/develop/DEVELOPMENT.md#running-tests-on-meteor-core>

¹¹<https://github.com/vercel/next.js/blob/canary/contributing/core/testing.md#writing-tests-for-nextjs>

CPython.¹² In this case, the recommendations are diverse, including how to write effective tests, deciding where to place the tests, choosing the appropriate types of tests to write, selecting a suitable testing library, and more

Test types: We analyze the three test types of the test pyramid: unit, integration, and end-to-end (e2e). Overall, 74% (115 in 156) of the test documentation mentions at least one test type. The most common is unit test, present in 71% of the test documentation. For example, the test documentation of Langchain mentions: “Unit tests: run on every pull request, so they should be fast and reliable [...] Unit tests cover modular logic that does not require calls to outside APIs. If you add new logic, please add a unit test”.¹³ Integration test is found in 20.5% of the cases. Project Langchain also mentions: “Integration tests cover logic that requires making calls to outside APIs [...] If you add support for a new external API, please add a new integration test.”. Lastly, we have e2e test, which is found in 15.5%. For example, the test documentation of Next.js states: “End-to-end (e2e) tests are run in complete isolation from the repository”.¹⁴

Test coverage: Test coverage measures the percentage of code that is covered (and uncovered) by tests [13] and is typically used to gauge the test effectiveness [14]. We find that 25.5% (40 in 156) of the test documentation mentions test coverage. For example, CPython has specific documentation dedicated to test coverage: “Increase test coverage [...] Ideally we would like to have 100% coverage, but any increase is a good one”.¹⁵ Likewise, scikit-learn also mentions a coverage threshold: “We expect code coverage of new features to be at least around 90%”.¹⁶ Project Materialize details what to do with code that is hard to cover: “Try and cover as many cases as you can, but don’t worry if there are some edge cases. You can add comments describing some problematic edge cases in TODOs so we know about them.”.¹⁷

Mocking: When creating tests, developers may find dependencies that make the test harder to implement. In this scenario, they can use mocks (test doubles) to emulate the dependencies’ behavior, contributing to making the test fast, isolated, and deterministic [15]–[19]. We detect that only 9.5% of the test documentation refers to test doubles. For example, project Vuex presents how to abstract API calls: “When testing actions, we usually need to do some level of mocking - for example, we can abstract the API calls into a service and mock that service inside our tests”.¹⁸ Project Pandas details what a unit test should not access: “A unit test should not access a public data set over the internet due to flakiness of network connections [...] To mock this interaction,

¹²<https://devguide.python.org/testing/run-write-tests/#writing>

¹³<https://python.langchain.com/docs/contributing/testing>

¹⁴<https://github.com/vercel/next.js/blob/canary/contributing/core/testing.md>

¹⁵<https://devguide.python.org/testing/coverage>

¹⁶<https://scikit-learn.org/stable/developers/index.html>

¹⁷<https://github.com/Dogfalo/materialize/blob/v1-dev/CONTRIBUTING.md>

¹⁸<https://vuex.vuejs.org/guide/testing.html>

use the `httpserver` fixture”.¹⁹ Likewise, the OpenHands test documentation details: “When you launch an integration test, mock responses are loaded and used to replace a real LLM’s response, so that we get deterministic and consistent behavior, and most importantly, without spending real money”.²⁰

Best practices & tips: Lastly, we looked for best practices and tips in software testing. We find that only 9% (14 in 156) of the test documentation provides this information. For example, project Next.js recommends creating bug-reproducing tests: “Best Practices: When applying a fix, ensure the test fails without the fix. This makes sure the test will properly catch regressions”.²¹ The test documentation of jQuery provides tips to improve test performance: “Test Suite Tips: During the process of writing your patch, you will run the test suite MANY times. You can speed up the process by narrowing the running test suite down to the module you are testing [...]”.²²

Finding 2: (i) Test documentation commonly explains how to run tests (83.5%), but less often provides guidance on how to write tests (37%). (ii) Test documentation frequently covers unit tests (71%), but rarely addresses integration (20.5%) and end-to-end tests (15.5%). (iii) Other key testing aspects are also less frequently discussed: test coverage (25.5%), mocking (9.5%), and best practices (9%).

IV. DISCUSSION

Current status of test documentation for contributors. We find that test documentation for contributors has diverse but unbalanced test content (*e.g.*, more how-to-run than how-to-write, more unit than e2e tests). Moreover, we detect that some content, such as e2e tests, test coverage, mocking, and best practices, is underrepresented in the test documentation. Without clear guidance, contributors must make their own decisions on critical aspects, such as when to create e2e tests, which mocking framework to use, what code to mock, and what coverage threshold to adopt. We recommend that project maintainers include this test information to inform contributors better and prevent issues caused by inadequately submitted contributions [1]. Also, researchers can propose novel tools to automatically detect contribution guidelines that lack certain test-related content (*e.g.*, mocking) to warn maintainers.

Guidelines to create test documentation for contributors. Our goal in this research is to provide guidelines to help create better test documentation for contributors. RQ2 provides an initial step in this direction. For example, for test coverage, we found many concrete recommendations for contributors about coverage increments, coverage thresholds, code that is

hard to cover, and more. In the mock context, we found recommendations of how to mock, what to mock, and what tool to use. Regarding the test types, we found definitions of what characterizes a specific type of test, when a test should be created, and when a test should be run. Therefore, we envision that a deep qualitative exploration of the actual test documentation for contributors can reveal best practices and guidelines that can be reused by open-source projects.

V. LIMITATIONS

This study provides an initial assessment of test documentation for contributors. The test documentation was detected manually, with the support of two authors of the paper. As detailed in Section II, when no test documentation was found for a given project, a second author conducted a follow-up search to double-check and minimize the risk of false negatives. Moreover, RQ2 explored the content of such documentation. However, we did not assess *what* recommendations are actually provided for contributors. We plan to deeply explore the recommendations in further studies.

VI. RELATED WORK

Multiple studies focus on better understanding the content of readme files and contribution guidelines [2], [4], [12], [20], [21]. Liu *et al.* assessed the structure of readme files in Java projects [20]. The authors found that the majority of readme files do not align with the GitHub guidelines, but repositories whose readme files follow the GitHub guidelines tend to be more popular. Wang *et al.* also studied the correlation between readme files and project popularity [21]. Prana *et al.* found that information about the “What” and “How” of a repository is common in readme files [12]. Elazhary *et al.* explored the content of contribution guidelines and compared their actual practices with the prescribed contribution guidelines [2]. The authors found that most projects diverge significantly from the expected process. Other research shows that such documentation may support new contributors [4]. Despite the various studies on contribution guidelines, there is a lack of research addressing software testing in such guidelines. Our research aims to fill this gap in the literature.

VII. CONCLUSION AND FURTHER STUDIES

We presented an empirical study to understand better how software testing is approached in contribution guidelines of 200 Python and JavaScript projects. We revealed the following insights: (1) 78% of the projects include some form of test documentation for contributors; (2) they focus more on how-to-run rather than how-to-write; (3) unit tests are commonly covered, while end-to-end tests receive less attention; and (4) discussions on test coverage and mocking are rarer.

In future work, we plan to *qualitatively* explore test documentation to catalog best practices and guidelines for contributors. In addition to the studied test content, we plan to address other aspects relevant to testing that may appear in contributor guidelines, such as testing techniques.

¹⁹https://pandas.pydata.org/docs/dev/development/contributing_codebase.html

²⁰<https://github.com/All-Hands-AI/OpenHands/blob/main/tests/integration/README.md>

²¹<https://github.com/vercel/next.js/blob/canary/contributing/core/testing.md#best-practices>

²²<https://github.com/jquery/jquery/blob/main/CONTRIBUTING.md#test-suite-tips>

ACKNOWLEDGMENTS

This research is supported by CAPES, CNPq, and FAPEMIG.

REFERENCES

- [1] Setting guidelines for repository contributors, <https://docs.github.com/en/communities/setting-up-your-project-for-healthy-contributions/setting-guidelines-for-repository-contributors>, October, 2024.
- [2] O. Elazhary, M.-A. Storey, N. Ernst, and A. Zaidman, “Do as I do, not as I say: Do contribution guidelines match the GitHub contribution process?” in *International Conference on Software Maintenance and Evolution (ICSME)*, 2019, pp. 286–290.
- [3] Open Source Guides, <https://opensource.guide>, October, 2024.
- [4] I. Steinmacher, M. A. G. Silva, M. A. Gerosa, and D. F. Redmiles, “A systematic literature review on the barriers faced by newcomers to open source software projects,” *Information and Software Technology*, vol. 59, pp. 67–85, 2015.
- [5] J. Tsay, L. Dabbish, and J. Herbsleb, “Influence of social and technical factors for evaluating contribution in github,” in *International Conference on Software Engineering*, 2014, pp. 356–366.
- [6] R. Pham, L. Singer, O. Liskin, F. Figueira Filho, and K. Schneider, “Creating a shared understanding of testing culture on a social coding site,” in *International Conference on Software Engineering (ICSE)*, 2013, pp. 112–121.
- [7] L. Dabbish, C. Stuart, J. Tsay, and J. Herbsleb, “Social coding in GitHub: transparency and collaboration in an open software repository,” in *Conference on Computer Supported Cooperative Work*, 2012, pp. 1277–1286.
- [8] Make a README, <https://www.makeareadme.com>, October, 2024.
- [9] O. Dabic, E. Aghajani, and G. Bavota, “Sampling projects in github for MSR studies,” in *International Conference on Mining Software Repositories, MSR 2021*, 2021, pp. 560–564.
- [10] H. Borges, A. Hora, and M. T. Valente, “Understanding the factors that impact the popularity of GitHub repositories,” in *International Conference on Software Maintenance and Evolution*, 2016, pp. 334–344.
- [11] H. Borges and M. T. Valente, “What’s in a GitHub star? understanding repository starring practices in a social coding platform,” *Journal of Systems and Software*, vol. 146, pp. 112–129, 2018.
- [12] G. A. A. Prana, C. Treude, F. Thung, T. Atapattu, and D. Lo, “Categorizing the content of GitHub readme files,” *Empirical Software Engineering*, vol. 24, pp. 1296–1327, 2019.
- [13] A. Hora, “Excluding code from test coverage: practices, motivations, and impact,” *Empirical Software Engineering*, vol. 28, no. 1, p. 16, 2023.
- [14] Coverage.py, <https://coverage.readthedocs.io>, October, 2024.
- [15] G. Meszaros, *xUnit test patterns: Refactoring test code*. Pearson Education, 2007.
- [16] G. Pereira and A. Hora, “Assessing mock classes: An empirical study,” in *International Conference on Software Maintenance and Evolution (ICSME)*, 2020, pp. 453–463.
- [17] D. Spadini, M. Aniche, M. Bruntink, and A. Bacchelli, “To mock or not to mock? an empirical study on mocking practices,” in *International Conference on Mining Software Repositories (MSR)*, 2017, pp. 402–412.
- [18] —, “Mock objects for testing java systems: Why and how developers use them, and how they evolve,” *Empirical Software Engineering*, vol. 24, pp. 1461–1498, 2019.
- [19] M. Fazzini, C. Choi, J. M. Copia, G. Lee, Y. Kakehi, A. Gorla, and A. Orso, “Use of test doubles in android testing: An in-depth investigation,” in *International Conference on Software Engineering*, 2022, pp. 2266–2278.
- [20] Y. Liu, E. Noei, and K. Lyons, “How readme files are structured in open source java projects,” *Information and Software Technology*, vol. 148, p. 106924, 2022.
- [21] T. Wang, S. Wang, and T.-H. P. Chen, “Study the correlation between the readme file of github projects and their popularity,” *Journal of Systems and Software*, vol. 205, p. 111806, 2023.