

Do Comments follow Commenting Conventions?

A Case Study in Java and Python

Pooja Rani*, Suada Abukar*, Nataliia Stulova*, Alexandre Bergel†, Oscar Nierstrasz*

*Software Composition Group, University of Bern, Bern, Switzerland

†Department of Computer Science (DCC), University of Chile, Santiago, Chile

✉ scg.unibe.ch/staff

Abstract—Assessing code comment quality is known to be a difficult problem. A number of coding style guidelines have been created with the aim to encourage writing of informative, readable, and consistent comments. However, it is not clear from the research to date which specific aspects of comments the guidelines cover (e.g., syntax, content, structure). Furthermore, the extent to which developers follow these guidelines while writing code comments is unknown.

We analyze various style guidelines in Java and Python and uncover that the majority of them address more the content aspect of the comments rather than syntax or formatting. However, when considering the different types of information developers embed in comments and the concerns they raise on various online platforms about the commenting practices, existing comment conventions are not yet specified clearly enough, nor do they adequately cover important concerns. We find that developers of both languages follow the writing style and content-related comment conventions more often than syntax and structure types of conventions. Our results highlight the mismatch between developer commenting practices and style guidelines, and provide several focal points for the design and improvement of comment quality checking tools.

Index Terms—Comment analysis, Software documentation, Coding Style Guidelines, Coding Standards

I. INTRODUCTION

Developers use several kinds of software documentation, including design documents, wikis, and code comments, to understand and maintain programs. Studies show that developers trust code comments more than other forms of documentation [1]. As code comments are usually written in a semi-structured manner using natural language sentences, and they are not checked by the compiler, developers have the freedom to write comments in various ways [2], [3], [4].

To encourage developers to write consistent, readable, and informative code comments, programming language communities and several large organizations, such as Google and Apache, provide coding style guidelines that also suggest comment-related conventions [5], [6], [7]. These conventions cover various aspects of comments, such as syntactic, stylistic, or content-related aspects. For instance, “Use 3rd person (descriptive), not 2nd person (prescriptive)” is an example of a stylistic comment convention for Java documentation comments [5]. However, to what extent these aspects are covered within different style guidelines and languages is not known. Therefore, we formulate the question: **RQ₁**: Which

type of comment conventions are suggested by various style guidelines?

As high-quality comments support developers in understanding and maintaining their programs, it is essential to ensure the adherence of their comments to the style guidelines to evaluate the overall comment quality. Rani *et al.* have investigated class comments of Smalltalk and their adherence to the commenting conventions provided by a default template [8]. They found that Smalltalk developers follow writing style and content-related comment conventions more than 50% of the time, but they use inconsistent structure and formatting of comment content. As Java and Python are among the most popular languages in use, several research works have focused on studying comments in Java and Python ([3], [4]), some especially focusing on class comments [9]. However, it remains largely unknown whether Java and Python developers adhere to the commenting conventions suggested by the style guidelines or not. To obtain this understanding, we formulate another research question: **RQ₂**: To what extent do developers follow commenting conventions in writing code comments in Java and Python?

Our initial results show that the majority of style guidelines propose more content-related conventions than other types of conventions, but compared to the different types of content developers actually embed in comments ([3], [4], [9]), and the concerns they raise on online platforms (e.g., StackOverflow or Quora) regarding comment conventions [10], it is clear that existing conventions are neither adequate, nor precise enough. On the other hand, these style guidelines often include conventions that are not relevant or applicable in many cases, leading developers to ignore them.

When the conventions are applicable, developers often follow the writing style and content conventions (80% of comments), but violate structure conventions in Java and Python class comments (nearly 30% of comments), confirming the previous results for Smalltalk by Rani *et al.* [8]. Although the project-specific guidelines provide very few additional class comment conventions, these conventions are followed more often compared to the conventions suggested by the standard guidelines both in Java and Python class comments. The data related to RQ₁ and RQ₂ is given in the replication package.¹

¹<https://doi.org/10.5281/zenodo.5296443>

Our results highlight the mismatch between conventions suggested by the style guidelines but not followed by developers in their projects and vice versa. Verifying this mismatch is currently not well-supported by tools or linters. The linters currently support very limited comment conventions [11], [12], e.g., they check the presence or absence of code comments, or their formatting with code but not their content. Our results indicate the need to conduct extensive studies on (i) which comment-related conventions linters provide, (ii) how well linters cover comment conventions from various style guidelines, and (iii) building tools and techniques to reduce the mismatch between developer commenting practices and style guidelines.

II. STUDY DESIGN

Data collection. Rani *et al.* have previously investigated the adherence of Smalltalk class comments to the conventions provided by a default template [8]. To verify their results across other languages, we analyze class comments in Java and Python. Rani *et al.* also investigated class comments of diverse Java and Python projects that vary in terms of size, domain, and contributors [9]. We use their dataset of Java and Python class comments for our analysis (RQ1 and RQ2). Table I shows the list of these projects in the column *Project* for the selected languages shown in the column *Language*.

a) *Comment conventions in Style Guidelines (RQ₁)*: The goal of this RQ is to investigate the type of comment conventions (or rules) various style guidelines suggest. The term *comment convention* refers to suggestions or rules about various aspects of comments, such as syntax, formatting, content, or writing style. First, we analyze the coding style guidelines of the projects (listed in Table I) which correspond to the Java and Python projects in the Rani *et al.* dataset [9]. Then we extract the comment-related rules from these guidelines. Each project might refer to project-specific guidelines in addition to the standard coding guidelines to customize its coding style. The standard guidelines are used as the baseline guidelines in the majority of the projects and are often provided by the programming language community such as Oracle, PEP, or organizations such as Google, Apache. In contrast, the project-specific guidelines scope to the project and extend, clarify, or conflict the standard guidelines such as Pandas.² Table I shows if a project supports project-specific commenting guidelines (✓) or not (×) in the column *Project guideline*, in addition to the standard guidelines the project mentions on its web page (listed in the column *Standard guideline*).

As within a style guideline, comment conventions can be scattered across multiple paragraphs, we scan all sentences and select those that mention any convention about comments. A rule can target various types of comments, such as class, method, package, or inline comments, or part of comments (e.g., summary, parameters). In case a sentence targets multiple comment types, we split the rule for each type. In total, we collected 600 comment-related rules. We organized all rules

TABLE I
OVERVIEW OF THE SELECTED PROJECTS AND THEIR STYLE GUIDELINES.

Language	Project	Project guideline	Standard guideline
Java	Eclipse	✓	Oracle
	Hadoop	✓	Oracle
	Vaadin	✓	Oracle
	Spark	✓	Oracle
	Guava	×	Google
	Guice	×	Google
Python	Django	✓	PEP8/257
	Requests	✓	PEP8/257
	Pipenv	×	PEP8/257
	Mailpile	×	PEP8/257
	Pandas	✓	Numpy
	iPython	✓	PEP8/257, Numpy
	Pytorch	✓	Google

into a taxonomy of five main categories: *Content*, *Structure*, *Formatting*, *Syntax*, and *Writing Style*. If a rule does not fit any of these categories, we put it into the *Other* category. The rationale behind the taxonomy is that the approaches evaluating comment quality can focus on a specific aspect of comments they want to evaluate and improve. Categories such as *Content* and *Writing Style*, are considered important by Rani *et al.* in their work on evaluating Smalltalk comments [8].

According to Rani *et al.*, the *Content* category contains the rules that describe which type of information the comment should contain while the *Writing Style* category contains natural-language specific rules, such as grammar, punctuations, and capitalization. We added three more categories, following their methodology, to cover other aspects of comments. The *Formatting* category deals with the rules related to indentation, blank lines, or spacing. It often complements *Structure* category conventions. The category *Structure* contains the rules about organizing the text, or location of the information in comments. For example, how the tags/sections/information should be ordered in the comments. The *Syntax* category focuses on the syntax to write a specific type of comments, for instance, which symbol to use to denote comments. Then, we analyze the frequency of these categories in the style guidelines to answer the RQ₁.

b) *Adherence of comments to conventions (RQ₂)*: The goal of this RQ is to verify whether or not developers follow the comment conventions, i.e., the rules identified in the previous RQ, in practice in their projects. Currently, there are no tools available that automatically check all comment types against all rule types, therefore, from each project we manually validate a sample of class comments against all class comment related rules extracted from its standard and project-specific guidelines shown in Table I (390 rules out of 600 rules). For the scope of this work, we focus on the rule types that require manual validation due to limited tool support i.e., all types of rules other than *Formatting*, thus, we validate 270 rules against the sample class comments.

We use the dataset by Rani *et al.* that provides sample class

²https://pandas.pydata.org/pandas-docs/stable/development/contributing_docstring.html

comments of the selected projects shown in Table I [9]. They selected a statistically significant sample of class comments from all class comments of each project with 95% confidence and 5% margin of error, resulting in a total of 700 class comments for both languages. In case a comment follows a particular rule, we label the rule as *followed*, otherwise as *not followed*. There are often cases where a rule is not applicable to the comment due to the unavailability of that information in the comment, e.g., the rules verifying syntax, content, or style of the version information in a class comment cannot be checked if the version information is not mentioned in the comment. For such cases, we label such rules as *not applicable* to the comment. We exclude a few rules for now that cannot be verified with the current dataset due to the abstract nature of a rule, the unavailability of the symbols that denotes the class comment, or code associated with the class, e.g., to verify the Oracle rule “for the `@deprecated` tag, suggest what item to use instead”, class comment alone is not enough and require code of the class to verify the replacement item. We plan to extend the dataset with more comment types and the required data.

We measure how many comments follow a particular rule and how many do not. One author labels the comments, and the second author reviews the labeled comments. In cases where they do not agree, the third author is consulted, and conflicts are resolved using the majority voting mechanism (Cohen’s kappa=0.80).

III. EARLY RESULTS

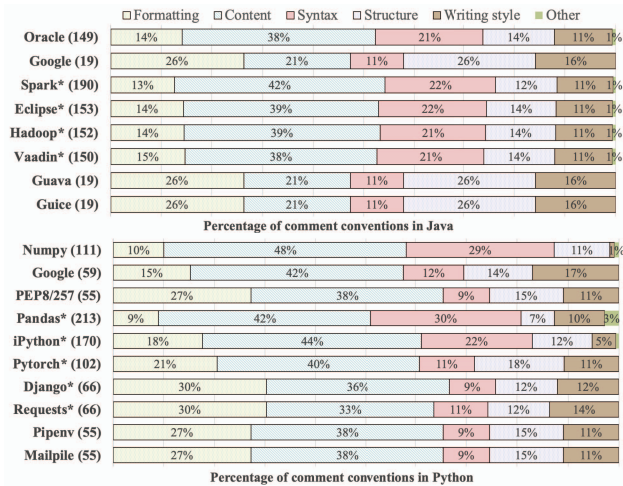


Fig. 1. Types of conventions in Java and Python guidelines

a) *Comment conventions in Style Guidelines (RQ₁)*: Figure 1 shows the total number of conventions for each standard guideline (Oracle and Google) and project-specific guidelines (including conventions from the standard guidelines) on the y-axis. The x-axis indicates the ratio of conventions belonging to a particular category from our taxonomy. Our results show that the majority of style guidelines present more rules about

the content to write (*Content*) in comments, except for the Google style guideline in Java, which contains more rules on how to format and structure comments (*Formatting*, *Structure*). Since the Oracle guideline is used as a baseline in several Java projects, project-specific guidelines suggest few additional comment conventions, and these conventions often either conflict with or clarify the standard guidelines. For example, the conventions, such as *line length limit* and *indentation with two spaces, four spaces, or tab*, are often among such additional and conflicting rules across projects. Identifying such rules and ensuring they are configured properly in tools can help developers in following them automatically.

Figure 1 also shows the distribution of rule types for Python style guidelines. Numpy and the projects following it (Pandas and iPython) contain the most rules about what type of information to write in comments and how to write it, compared to other standard guidelines, such as those of Google, PEP, and Oracle. For example, the Numpy guideline suggests writing a short and extended summary of the class, usage examples, notes and warnings in a class comment, and provides syntax and style conventions to write these types of information. Class comments of iPython and Pandas contain all of these information types and follow the syntax conventions to write them. Interestingly, developers write such types of information in all other projects [4], [9] regardless of whether the project guideline suggest or not, but they are writing these information types in inconsistent ways. Previous comment analysis studies for Java and Python also show that developers embed other, different types of information in comments, such as *Usage*, *Expand*, *Rationale*, or *Pointer*, but we do not find conventions in the corresponding style guidelines (Google, PEP, Oracle) to write such types of information [3], [4].

We observe that even though style guidelines are intended to encourage and help developers to write good comments for all code entities, comment conventions are scattered across multiple sources, documents, and paragraphs. Thus, it is not always easy to locate conventions particular to one entity (class, function, inline), causing developers to seek conventions using online sources [10].

Finding. The majority of the style guidelines propose more content-related conventions than other types of conventions, but they are not easy to locate in the style guidelines, and do not always match developer commenting practices.

Finding. The Numpy style guideline provides more rigorous content conventions for comments compared to other style guidelines, such as Oracle, PEP257, or Google.

b) *Comment conventions in Style Guidelines (RQ₂)*: Figure 2 shows the distribution of comments within each project that follows rules, do not follow them, or to which the rules are not applicable. For example, in Eclipse on average 27% of the comments follow the rules, whereas 3% of the comments violate the rules, and 70% of the comments do not have enough relevant information within them to check them against a rule. High ratio of non applicable rules (shown in Figure 2) to




Fig. 2. Percentage of comments that follow rules, do not follow them, or to which rules are not applicable.

selected comments indicates that the style guidelines suggest various comment conventions, but developers rarely adopt them while writing comments. For instance, the Oracle rules in Java, such as “use *FIXME* to flag something that is bogus or broken”, or “use *@serial* tag in class comment” are not applicable on comments due to unavailability of *FIXME* or *@serial* information in comments and thus showing the developers interest in not adopting such rules. Similarly, some rules in Python, such as “*Docstrings for a class should list public methods and instance variables*” are also rarely adopted.

Finding. Style guidelines suggest various comment conventions, but developers do not or rarely adopt them while writing comments.

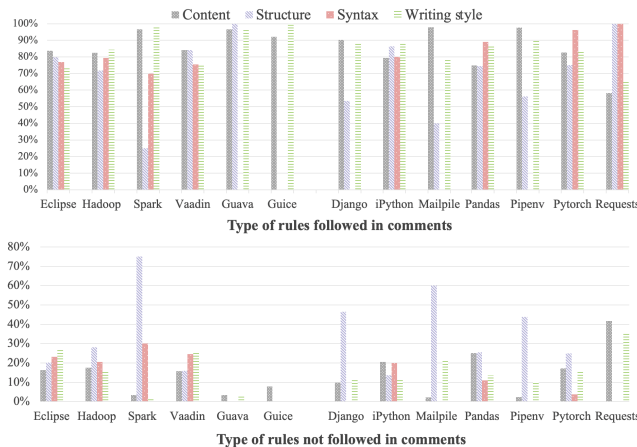


Fig. 3. Types of rules followed or not in Java and Python projects

Of the rules that are applicable as shown in Figure 3, writing style and content rules were more often followed than syntax and structure rules, confirming previous results for Smalltalk [8]. It shows that developers are interested in writing informative and consistent comments.

Finding. Compared to Python, Java class comments violate rules less often (as shown in Figure 3).

Finding. Class comments in Java and Python often follow writing style and content conventions (80% of comments), but violate structure conventions (30% of comments).

As discussed before, some rules are often followed, while there are others that are frequently violated. For instance, the syntax rule “*separate the paragraphs with a `<p>` paragraph tag*” in Spark is violated often. Similarly, in Pandas the rule “*a few sentences giving an extended summary of the class or method after the short (one-line) summary*” is often followed but the rule “*there should be a blank line between the short summary and extended summary*” is often violated. Such conventions can be further investigated by surveying developers to know the specific factors, such as the usage of linters for comments, team strictness, or developer awareness behind these explicit instances of rule adherence or violation. Although the project-specific guidelines, as shown in Table I, provide few additional conventions, these conventions are followed more often in the projects compared to the conventions provided by their standard guidelines. Specifically, 85% of Python class comments and 89% of Java class comments follow the project-specific conventions, whereas 81% of their comments follow the conventions from the standard guidelines. One such example is, the rule “*Do not use @author tags*” is specific to Hadoop and in contrast with the Oracle style guideline, but it is always followed in Hadoop comments. It would be an interesting future work to explore the reasons behind such conventions.

Finding. Project-specific class comment conventions are followed more often than the conventions suggested by the standard guidelines.

IV. IMPLICATION & RELATED WORK

Impact of commenting conventions. Coding style guidelines impact program comprehension and maintenance activities. However, not all conventions from the guidelines have the same impact on these activities. Smit *et al.* [11] ranked 71 code conventions that are most important to maintainable code. However, they accounted only for missing or incomplete Javadoc comments on public types and methods, and did not account for other comment-related conventions, especially about their content. Similarly, most previous work has focused on building tools for formatting and naming conventions for code entities, while being very limited on comment conventions [2], [13]. We provide a dataset of 700 labeled class comments and 600 comment conventions (taxonomy) for Java and Python. This dataset can help researchers rank the specific comment conventions to find out their importance, and impact on the program comprehension and maintenance activities, and thus help in developing comment quality tools based on the supervised machine-learning approaches.

Comment generation. To reduce developer effort in writing comments, various researchers have proposed to generate

comments automatically. Moreno *et al.* proposed a template-based approach to generate class comments in Java [14]. Given the importance of including developer commenting practices in such approaches, and the impact of a template on developers [8], our results can help researchers design such templates more carefully.

Adherence of comment conventions. Previous works, including Bafatakis *et al.* and Simmons *et al.*, evaluated the compliance of Python code to Python style guidelines [15], [12]. However, they included only few comment conventions and missed many other content and writing style conventions. In our study, we find various comment conventions, such as grammar rules, the syntax of writing different types of information that developers often follow, but which are not covered in such studies. Rani *et al.* measured the adherence of Smalltalk class comments to the default comment template and found that developers follow the writing conventions of the template [8]. Java and Python do not provide any default template to write comments but support multiple style guidelines for each project, thus collecting and verifying their comment conventions against comments is more tricky. We study diverse projects in Java and Python and found that developers follow writing and content conventions more than other types, thus confirming the results of Rani *et al.* for Smalltalk.

V. CONCLUSIONS

Given the importance of code comments and consistency concerns in projects, we study various style guidelines and diverse open-source projects in the context of comment conventions. We highlight the mismatch between what conventions the style guidelines suggest for class comments, and how often developers adopt and follow them, and what conventions developers follow in their class comments but which are not suggested or mentioned by the style guidelines. However, identifying automatically this mismatch is not yet fully achieved. This indicates the need to further automate the software documentation field. Our results also indicate the need to conduct extensive studies on various linters or quality assessment tools to know the extent they cover various comment conventions and improve them for the missing conventions. In this direction, we provide a dataset of labeled conventions against comments, and the methodology to extract conventions from the guidelines and verify them against comments. This can help in expanding the scope of the work. We plan to expand the study to other types of comments, conventions, and programming languages to generalize the current results and design the tools accordingly. Having the tools to automatically assess the documentation quality at each development stage can help developers in maintaining overall high-quality documentation.

VI. ACKNOWLEDGEMENT

We gratefully acknowledge the financial support of the Swiss National Science Foundation for the project “Agile Software Assistance” (SNSF project No. 200020-181973, Feb 1, 2019 - Apr 30, 2022). Bergel is also grateful to Lam Research and the

ANID FONDECYT Regular 1200067 for partially sponsoring the work presented in this paper.

REFERENCES

- [1] W. Maalej, R. Tiarks, T. Roehm, and R. Koschke, “On the comprehension of program comprehension,” *ACM TOSEM*, vol. 23, no. 4, pp. 31:1–31:37, Sep. 2014. [Online]. Available: <http://mobis.informatik.uni-hamburg.de/wp-content/uploads/2014/06/TOSEM-Maalej-Comprehension-PrePrint2.pdf>
- [2] M. Allamanis, E. T. Barr, C. Bird, and C. Sutton, “Learning natural coding conventions,” in *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. FSE 2014. New York, NY, USA: ACM, 2014, pp. 281–293.
- [3] L. Pascarella and A. Bacchelli, “Classifying code comments in Java open-source software systems,” in *Proceedings of the 14th International Conference on Mining Software Repositories*, ser. MSR ’17. IEEE Press, 2017, pp. 227–237.
- [4] J. Zhang, L. Xu, and Y. Li, “Classifying python code comments based on supervised learning,” in *Web Information Systems and Applications - 15th International Conference, WISA 2018, Taiyuan, China, September 14-15, 2018, Proceedings*, ser. Lecture Notes in Computer Science, X. Meng, R. Li, K. Wang, B. Niu, X. Wang, and G. Zhao, Eds., vol. 11242. Springer, 2018, pp. 39–47.
- [5] “Oracle documentation guideline,” verified in Aug. 2021. [Online]. <https://www.oracle.com/technetwork/java/javase/documentation>
- [6] “Python documenation guideline,” verified in Aug. 2021. [Online]. <https://www.python.org/doc/>
- [7] “Google style guidelines,” verified in Aug. 2021. [Online]. <https://google.github.io/styleguide/>
- [8] P. Rani, S. Panichella, M. Leuenberger, M. Ghafari, and O. Nierstrasz, “What do class comments tell us? An investigation of comment evolution and practices in Pharo Smalltalk,” *Empirical Software Engineering*, vol. 26, no. 6, pp. 1–49, 2021.
- [9] P. Rani, S. Panichella, M. Leuenberger, A. Di Sorbo, and O. Nierstrasz, “How to identify class comment types? A multi-language approach for class comment classification,” *Journal of Systems and Software*, vol. 181, p. 111047, 2021.
- [10] P. Rani, M. Birrer, S. Panichella, M. Ghafari, and O. Nierstrasz, “What do developers discuss about code comments?” in *2021 IEEE 21st International Working Conference on Source Code Analysis and Manipulation (SCAM)*, 2021.
- [11] M. Smit, B. Gergel, H. J. Hoover, and E. Stroulia, “Maintainability and source code conventions: An analysis of open source projects,” *University of Alberta, Department of Computing Science, Tech. Rep. TR11*, vol. 6, 2011.
- [12] A. J. Simmons, S. Barnett, J. Rivera-Villicana, A. Bajaj, and R. Vasa, “A large-scale comparative analysis of coding standard conformance in open-source data science projects,” in *Proceedings of the 14th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, 2020, pp. 1–11.
- [13] M. Arai, “Development and evaluation of eclipse plugin tool for learning programming style of java,” in *2014 9th International Conference on Computer Science & Education*. IEEE, 2014, pp. 495–499.
- [14] L. Moreno, J. Aponte, G. Sridhara, A. Marcus, L. L. Pollock, and K. Vijay-Shanker, “Automatic generation of natural language summaries for Java classes,” in *IEEE 21st International Conference on Program Comprehension, ICPC 2013, San Francisco, CA, USA, 20-21 May, 2013*, 2013, pp. 23–32.
- [15] N. Bafatakis, N. Boecker, W. Boon, M. C. Salazar, J. Krinke, G. Oznacar, and R. White, “Python coding style compliance on stack overflow,” in *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 2019, pp. 210–214.