

Analyzing the Quality of Submissions in Online Programming Courses

Maria Tigina
JetBrains Research
Republic of Serbia
maria.tigina@jetbrains.com

Anastasiia Birillo
JetBrains Research
Republic of Serbia
anastasia.birillo@jetbrains.com

Yaroslav Golubev
JetBrains Research
Republic of Serbia
yaroslav.golubev@jetbrains.com

Hieke Keuning
Utrecht University
The Netherlands
h.w.keuning@uu.nl

Nikolay Vyahhi
Stepik
United States
vyahhi@stepik.org

Timofey Bryksin
JetBrains Research
Republic of Cyprus
timofey.bryksin@jetbrains.com

Abstract—Programming education should aim to provide students with a broad range of skills that they will later use while developing software. An important aspect in this is their ability to write code that is not only correct but also of high quality. Unfortunately, this is difficult to control in the setting of a massive open online course. In this paper, we carry out an analysis of the code quality of submissions from JetBrains Academy — a platform for studying programming in an industry-like project-based setting with an embedded code quality assessment tool called Hyperstyle. We analyzed more than a million Java submissions and more than 1.3 million Python submissions, studied the most prevalent types of code quality issues and the dynamics of how students fix them. We provide several case studies of different issues, as well as an analysis of why certain issues remain unfixed even after several attempts. Also, we studied abnormally long sequences of submissions, in which students attempted to fix code quality issues after passing the task. Our results point the way towards the improvement of online courses, such as making sure that the task itself does not incentivize students to write code poorly.

Index Terms—programming education, code quality, MOOC, learning programming, refactoring, large-scale analysis

I. INTRODUCTION

Nowadays, software lies at the heart of virtually every area of our life [1], including such crucial fields as medicine, banking, and governance, where the cost of even a single error can be incredibly high [2], [3]. Since the quality of code directly affects the maintainability, flexibility, and performance of the software [4]–[6], maintaining its high quality also becomes a priority [7]. However, many novice programmers do not pay enough attention to the quality of the software they develop [8], which indicates the importance of instilling the ability to write high-quality code during the education process [9], [10].

A popular way of learning programming is attending massive open online courses (MOOCs) [11], which answer the needs of a growing number of students, in particular during the pandemic [12]. While MOOCs allow everyone to get the programming education they require [13], they provide significantly less control over the student and their progress [14],

which means that the student may pass the course without learning everything they need to know. Specifically, it is possible to submit solutions to programming tasks without maintaining the desired level of code quality, because, on the one hand, it is not feasible to manually check this many submissions [15], and on the other, professional automated code quality tools are not adapted to the education process [10].

Much research has been conducted that studies the aspect of code quality in code written by students in MOOCs or regular classes [8], [16]–[21]. The majority of the existing research focuses on Java [8], [16]–[18] and Scratch [19], [20], since public large-scale datasets already exist for these languages [22], [23]. Despite the growing popularity of Python in education [24], Python code quality studies are usually limited to submissions from just one or several semesters of a university course [18], [25]. Moreover, researchers usually only consider code fragments of tasks for beginners [8], [18]–[20], whereas it would also be valuable to look at more complex code snippets. Finally, existing works do not take into account external factors of code quality issues, *e.g.*, the issues in the task itself, which can affect the results.

To bridge the existing gaps in research, in this work, we conduct a large-scale analysis of Java and Python student submissions from a popular MOOC platform — JetBrains Academy [26]. The platform’s team provided us with a large dataset of 1,073,018 Java submissions and 1,345,332 Python submissions for more than 700 different tasks. This includes only the correct submissions, after passing all tests, allowing us to study code quality issues in them. We applied a tool called Hyperstyle [27] to discover code quality issues in the submissions. Hyperstyle is embedded into JetBrains Academy, meaning that students see its feedback after a successful task submission and can fix detected code quality issues in their next attempts to increase their code quality grade. This allowed us to use the historical data of past submissions to analyze the way students resolve existing issues.

The analysis of the most popular code quality issues showed that among both languages, not only minor issues like unused

imports are popular, but also issues that can lead to bugs in the future, *e.g.*, the incorrect usage of a switch statement in Java or the shadowing of built-in function names in Python. This indicates the necessity to teach students about them early on. Moreover, we found that frequent code quality issues are not always the fault of students, and can be caused by tasks on the platform itself: for example, incomplete or incorrect theory part, a weak test base, mistakes in task descriptions or in pre-written code templates, an incorrect order of topics that leads to the misunderstanding of the learned concepts, etc.

Next, we looked at the way students fix their issues. We found out that the majority of students are good at fixing simple issues that do not require big changes in the code. Moreover, some students even tried to correct code quality issues that were not their fault, *e.g.*, fix issues in a pre-written template or in a code sample taken from an incorrect theory part. This indicates a successful synthesis of the code quality analyzer and the platform, since students try to follow the advice of the code quality analyzer to fix issues even if they did not introduce them. During the detailed manual analysis of the submissions, we discovered several interesting cases, *e.g.*, many students try to shorten their code and only worsen its quality due to the lack of experience in refactoring.

Finally, we analyzed abnormally long sequences, where after successfully passing the task, students submit five or more attempts, changing or fixing the code. It turned out that students can spend dozens of attempts improving their code and get the highest code quality score. We noticed that the main reason for many attempts is the student's lack of understanding of what exactly needs to be corrected. For example, in the code editor on the platform, only the line with the issue is highlighted, but not the exact position, which confuses students, so they try different fixes until they reach the correct code. Another interesting observation is that students correct issues by groups, *i.e.*, if students see a familiar issue and know how to correct it, they try to correct it in all the places where it appears at once.

The results of our study and the provided insights can be of use to various groups of practitioners. Teachers who are developing or maintaining programming courses can use our results to modify their content to focus on the most prevalent issues. Software developers who are creating new code quality tools can consider our insights to take care of the most prominent mistakes, as well as to facilitate their better fixing among students. Finally, content managers and creators of MOOCs and their tasks can ensure that the content itself does not facilitate the issues, be that in the theory or templates.

The tool that we developed for gathering data from JetBrains Academy is available online, among with the detailed statistics about the code quality issues within our dataset [28], so others can extend this paper.

Overall, the main contributions of our paper are:

- 1) **Analysis** of more than a million successful submissions in Java and more than 1.3 million successful submissions in Python from the standpoint of their quality. We describe

the most popular types of code quality issues and analyze the dynamics of students fixing them.

- 2) **Case studies** of the most prominent issues, where we provide several examples that demonstrate how students struggle with code quality problems and how a MOOC's environment can help or interfere with the studying process.
- 3) **Insights** and practical implications that can be useful to teachers who are in the process of developing courses, as well as developers working on MOOC platforms and their quality control tools.

II. BACKGROUND

A. Code Quality Analysis

Many works study the problems of students' code quality, analyzing large-scale datasets of student submissions to find certain common behaviors and patterns [8], [17]–[20], [29] and testing code quality tools on small groups of students [25], [30], [31]. The first group mostly focuses on Scratch and Java languages, and uses open datasets for their research [22], [23].

Keuning et al. [8] studied Java submissions from the Blackbox database [22]. The authors used PMD [32], a pre-configured industrial code analyzer, to find common code quality issues. In addition, the authors studied which issues students correct over time, and compared code quality of students who use code analyzers with those who do not. The main weak point of this work is the lack of studying issues of more experienced students [17].

Edwards et al. [17] broadened the research towards more difficult tasks. They studied over 500,000 Java submissions from students with different programming experience, and categorized the obtained code quality issues, especially of formatting issues. The students used the Web-CAT tool [33] to submit their solutions and immediately receive feedback from professional analyzers. In addition to researching general issues, the authors explored the issues with more granular metrics, for example, which issues take more time to fix. The authors found that students correct issues poorly since professional linters are not adapted to the educational process, *i.e.*, provide feedback that is difficult for students to understand [10], [17].

Albluwi et al. [18] studied the distribution of code quality issues from professional linters. The authors searched for a correlation between the issues and various metrics, such as student experience and progress in the course, course difficulty, etc., but the size of the dataset they used was rather small. The data includes submissions of 968 students from three semesters of an introductory programming course that solve nine programming tasks, which might be too small to generalize the results. The above-mentioned problem of having unadapted feedback is also present.

The next group of works focused on Python. The PYTA tool was developed by Liu and Petersen [25] to provide custom checks for common code quality issues with adapted feedback. This tool is a wrapper for the popular Pylint [34] analyzer. The authors compared the issues before and after embedding the

tool into a computer science course, and found that students who use the tool make fewer mistakes and spend less time fixing them. The data consists of 114,865 submissions from 40 coding exercises, however, there is no detailed analysis of the code quality issues in these solutions, and the work is aimed only at the beginner students.

Molnar et al. [30] analyzed Python code in submissions of various difficulty using Pylint [34] and employed their own tool for visualizing the issues. The authors used a special score computed by Pylint to assess the quality of the code. The work shows how the code quality score changes depending on the complexity of the tasks, as well as depending on the progress of the given student. However, the dataset only contains 642 submissions, and, similarly to previous works, an unadapted code quality tool was used for the assessment.

Effenberger and Pelánek [29] recently analyzed code quality issues in 114,000 Python submissions for 161 tasks, taken from 11,000 students over 2.5 years. The authors used the data for tasks that can be solved with 2–20 lines of code with a median time for their solving from 1 to 20 minutes. The authors analyzed not only the issues that exist in correct solutions, but also how the provided feedback affects their correction. However, the main limitation of this work is the simplicity of the tasks.

Overall, it can be seen that the topic of the quality of student code is important and rather well-researched, however, existing studies share a number of common drawbacks. The open-access datasets (e.g., Blackbox [22]) are large, but they do not contain detailed information about the student’s full context, like the particular task the student is working on [8]. Other datasets are significantly smaller—from several hundred [30] to a hundred thousand submissions [25]—and only contain tasks for novices. Finally, researchers use tools, the output of which is not adapted to the educational process, and students may find it difficult to interpret [10].

In our work, we aim to overcome these limitations by conducting a study using a large dataset of solutions from a popular education platform JetBrains Academy [26] that employs a code quality tool called Hyperstyle [27] within it.

B. JetBrains Academy and Hyperstyle

JetBrains Academy [26] is a project-based education platform developed by JetBrains [35]. The learning on the platform is organized as completing a chain of small tasks that are all part of a single complex one, thus, after solving them, the student has a finished project. Currently, the platform supports studying Java, Python, Kotlin, JavaScript, and Go. The platform can be used by students, but is also available for everybody. Some courses are free, some require a paid subscription, both Java and Python courses that are the focus of this study are paid.

JetBrains Academy is built around a large knowledge graph, reflecting the set of topics necessary for studying a programming language. The topics can be both theoretical (e.g., asymptotic complexity) and practical (e.g., arrays) with some topics being common between languages (e.g., algorithms

or data structures). To learn more complex concepts, it is necessary to move in deeper along the graph. Each task in a studied project is connected to one topic in the graph, aligning the development of the project with the learning of concepts. This allows defining the *complexity* of the task, i.e., the depth in the knowledge graph of the topic that corresponds to this task. There are three levels of complexity depending on the depth: *shallow* (≤ 2), *deep* (≥ 5), otherwise *moderate*.

Submissions to programming tasks are validated in two stages. Firstly, the correctness of the solution is checked using tests, similar to most MOOCs. If all the tests pass, the second stage is initiated where code quality is checked using the Hyperstyle [27] tool. Hyperstyle uses a subset of code quality checks of professional analyzers (like PMD [32] for Java or Pylint [34] for Python) that are relevant to the educational process. Also, it changes the messages for the majority of the analyzers’ issues to be more understandable for students. As a result, students receive a grade based on the overall quality of their code. Different code quality issues affect the final grade differently, some issues are only informative (*minor* issues) and do not lower the grade at all.

The detected issues are also highlighted in the JetBrains Academy platform’s user interface, so that the students can see and fix them in subsequent attempts and improve their grades. The UI of JetBrains Academy, together with the Hyperstyle issue and output, are presented in Figure 1. It is also important to mention that students who have successfully passed the task can see other public submissions of fellow students with highlighted code quality issues.

All of this makes JetBrains Academy a suitable platform to study code quality issues. It already comes with the embedded code quality tool, which allows us not only to analyze the issues themselves, but also check the history of submissions from students to track how they fixed these issues.

III. DATASET

For our analysis, the JetBrains Academy team provided the data about students’ successful submissions on the platform over one year (from September 1st, 2020 to September 1st, 2021) that are written in Java or Python and passed all the correctness tests for the corresponding task. Every submission contains the source code of the solution, various metadata like timestamp, task ID, user ID, etc., as well as platform-specific features such as the task’s complexity. All personal information about users was completely removed by the platform’s team. The obtained data was preprocessed and then filtered.

Preprocessing. All successful solutions submitted by each user to a particular task consecutively were grouped and sorted by timestamp, we will refer to such a collection as a *submission series*. If two consecutive submissions in one series were identical, the latter was filtered out. Next, for each submission, we ran the Hyperstyle tool and collected its output to see what code quality issues are present in it. We use the output of the Hyperstyle tool, since this tool is used on the platform to show code quality issues to the students.

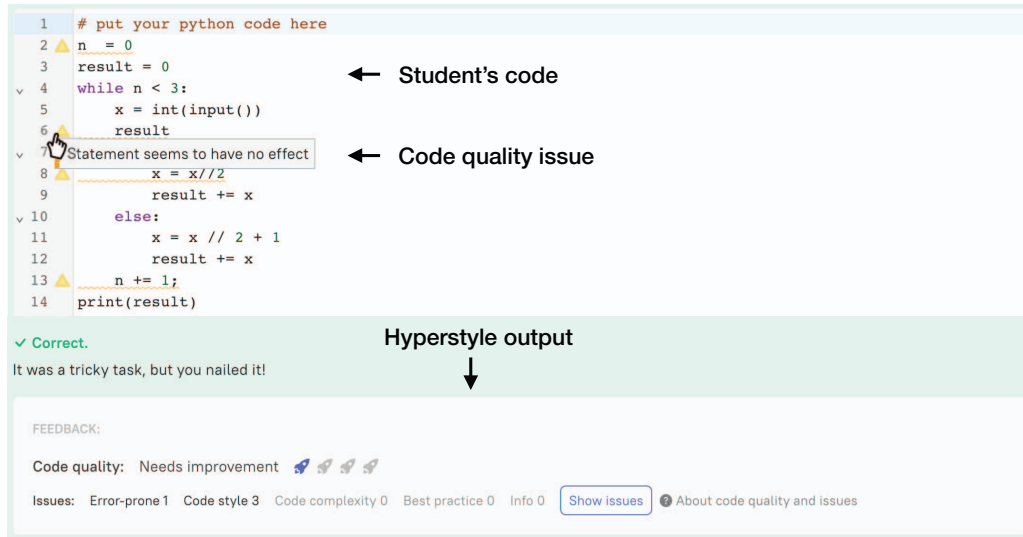


Fig. 1. The UI of a task on JetBrains Academy, together with a Hyperstyle issue tooltip in the editor and the tool's output.

TABLE I
CHARACTERISTICS OF THE COLLECTED DATASET
AFTER PREPROCESSING AND FILTERING.

	Java	Python
Tasks	415	364
Students	37,892	46,863
Solutions to shallow tasks	292,049	380,712
Solutions to moderate tasks	428,110	618,257
Solutions to deep tasks	352,859	346,363
Total solutions	1,073,018	1,345,332

Filtering. Firstly, we omitted issues found in *templates* — pre-written parts of the solution from the creators that students can edit and supplement with their own code. Issues in the template were introduced by the task creators, not students, so we did not take them into account in our analysis. To filter them, we applied an algorithm that is used on the platform itself to find code quality issues in the templates and fix them. The algorithm analyzes students' solutions, looks for issues that are most frequently not corrected, matches positions with the template, and finally marks them as template issues.

In order to select more representative and interesting data for the general analysis, we applied a number of filtrations. Firstly, for the main analysis we kept only the submission series that contain five successful submissions or fewer. This allowed us to remove suspiciously long series but filter out only 5% of the data, since the vast majority of users finish the task in under five attempts. Secondly, we skipped all formatting issues like incorrect indentation or missing whitespaces to focus on more prominent and serious issues. Although formatting issues are important and appear in student's solution more often, they are already well-studied [17], [18], [25] and the process of correcting them is rather straightforward. The characteristics of the final dataset are presented in Table I.

IV. ANALYSIS

In our analysis, we studied the following two research questions about the code quality of student submissions.

RQ1: Which code quality issues are the most prevalent in Java and in Python submissions?

RQ2: How do students fix various types of issues as they update their submissions with further attempts?

A. RQ1: Most Prevalent Issues

In this section, we study the most frequent issues and the reasons for their popularity.

Methodology. In each submission series, we selected only the first attempt, since we wanted to focus on the students' initial issues, before they had seen any feedback from the linters. These attempts represented 78% of all submissions for Java and 86% for Python, respectively. For each issue type, we calculated the percentage of submissions in which it was detected at least once, separately for the tasks with different complexity levels. Finally, we compared the most prevalent issues and the differences between complexities.

Results. Firstly, it is important to note that only 8% of Java and 11% of Python first attempt submissions have at least one code quality issue not connected to formatting. The percentages for each issue type are presented in Figure 2, the top part shows the most popular issues in Java, while the bottom part shows the information for Python. Both parts are sorted by the total percentage among all submissions.

The Top-5 Java issues include: (1) *Unused local variable*, which finds a locally declared variable that is not used; (2) *Unused imports*, which is similar in regards to imports; (3) *Missing break in switch*, which points to a possibly skipped `break` operator in a casing construction; (4) *Boolean expression complexity*, which warns against conditions that are too complex, e.g., `if (a > -15 && a <= 12 || a > 14`

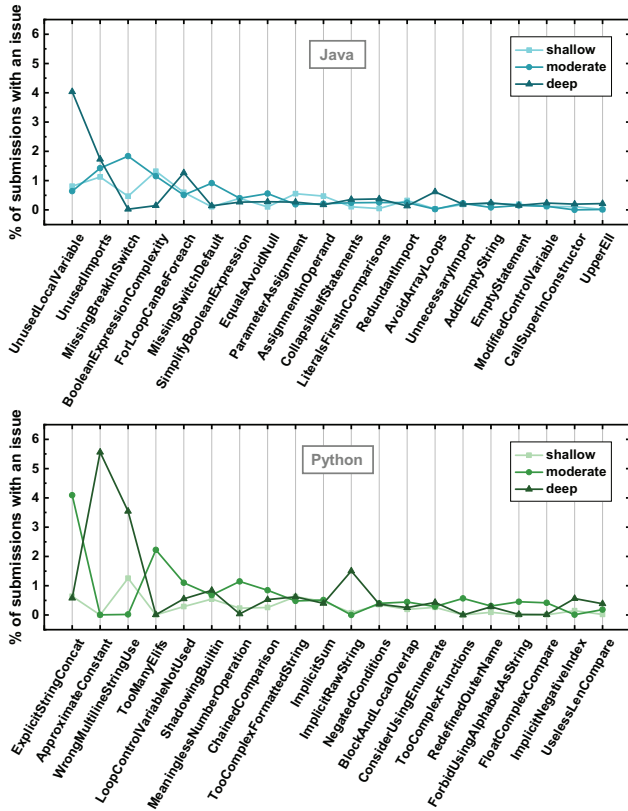


Fig. 2. Distribution of the most prevalent issues among Java submissions (top) and Python submissions (bottom). The issues are sorted by their total prevalence in all submissions.

`&& a < 17 || a >= 19) { . . . };` and (5) *For loop can be foreach*, which advises to rewrite the `for` loop into the `foreach` to avoid explicit indexing.

The Top-5 Python issues include: (1) *Explicit string concatenation*, which forbids using string concatenation; (2) *Approximate constant*, which finds the constants that the students defined and that can be replaced by existing constants from the `math` library; (3) *Wrong multiline string use*, which informs about the incorrect usage of multiline strings as function arguments; (4) *Too many elifs*, which informs about using `elif` statements that are too complex; and (5) *Loop control variable not used*, which finds unused variables in loops that should be replaced with the underscore symbol.

While there exists a difference between different task complexities, we cannot definitely conclude that some issues are more popular in moderate tasks than in deep ones. Usually, the high frequency of the issue can be explained by the task itself. For example, *Missing break in switch* is more popular in moderate tasks, because students learn this construct in a task with such complexity. The same findings were highlighted by Effenberger and Pelánek [29]. They found that some issues are highly localized in several tasks and thus influence the general distribution. In Section V, we touch upon some of them in more detail and give concrete code examples to illustrate the most peculiar situations.

B. RQ2: Dynamics of the Solution Quality

In this section, we study how students fix quality issues in their submission series. While the Hyperstyle paper [27] has studied the influence of the tool on the correction of code quality issues, a detailed analysis of which issues get corrected was not provided.

Methodology. Firstly, we filter out all submissions with only a single attempt, which is 81% of all submission series for Java and 87% for Python. Of these, less than 10% have at least one code quality issue, however, the students chose not to fix them. Next, for each submission series that we kept, we calculate the percentage of issues left in the last submission in the series relative to the first attempt. For example, if there were 1,000 initial submissions with the given issue, and 200 in the last attempts, we say that the issue remained in 20% of cases. Comparing the first and last attempts allows us to focus on the high-level, overall nature of fixing issues.

Results. The dynamics of fixing various issues in submission series is shown in Figure 3, the top chart represents Java submissions, while the bottom one represents Python submissions. In the chart, we consider the most prevalent issues that were identified in RQ1.

For Java, many issues are successfully corrected by students and remain in the final submission in less than 15% of cases: (1) *Upper ell* forbids the student to use a lower-case `l` in numbers of the `long` type, since the lower-case `l` looks a lot like `1`; (2) *Empty statement* informs the student about standalone “;” semicolon with no statement; and (3) *Unused imports*. These issues are easy to recognize and understand. Moreover, to correct such issues, students need to add or delete a single line or keyword, which explains the fact that students successfully fix them. However, several issues are not corrected in more than 90% of cases: (1) *Call super in constructor* reminds of the need to call the constructor of the class of the child; (2) *Add empty string* indicates a useless concatenation with an empty string. To make matters worse, there is a 14% increase in submissions with the *Assignment in operand* issue, which prohibits an assignment within the conditional expressions of `while` or `if` statements. We examine the reasons for this anomaly in detail in Section V-B. Keuning et al. [8] also showed that some mistakes are not corrected by students, while some are corrected by the majority.

For Python, there are only two issues which remain in less than 15% of students’ submissions: (1) *Shadowing Python built-in*, which appears if students override built-in functions like `max` or `input`; and (2) *Loop control variable not used* when loop control variable is not in use and can be replaced with an underscore. Moreover, unlike Java, most of the issues in the studied list are not corrected by students even in the final attempt in more than 50% of cases.

Finally, as many as three of them appear more often in the last attempt than in the first: (1) *Too complex formatted string*, which is connected to using `f-strings` incorrectly (strings like `f"My name is: {name}"`); (2) *Wrong multiline string use*, which forbids using multiline strings directly;

The percentage of issues left in the last submission relative to the first submission




Fig. 3. The percentage of submission series that have each issue left in the last attempt, for Java (top) and Python (bottom). Lower values indicate better fixing, higher values indicate worse fixing.

(3) *Implicit raw string* which forbids escape sequences inside regular strings and forces to use Python raw strings instead. Similar to Java, we analyzed these anomalies separately in the case studies in Section V-B, and came to the conclusion that in all cases it is not the fault of the students. The reason lies in the platform itself, namely, in the incomplete or erroneous theoretical part or the task idea and formulation in general, which lead to the inability to implement the task correctly.

V. CASE STUDIES

Generally, it is difficult to find the true root cause of the popularity of an issue without manually studying specific cases. In this section, we describe several cases that provide explanations for some major peaks in our data.

```

a) 1 import java.util.Arrays;
    2 import java.util.Scanner;
    3
    4 public class Main {
    5     public static int maxAbs(String[] num){
    6         // Expected solution
    7         Arrays.stream(num)...
    8
    9         // Student solution#1
   10        Stream.of(num)...
   11
   12        // Student solution#2
   13        for (int i = 0; i < num.length; i++)
   14            {...}
   15    }
   16}

The unused import.
Students use an alternative
strategy to solve the task
(lines 9-13).

"Equals" skips the null check.
It is better to use:
if ("0".equals(input)) {...}

b) 1 while (scanner.hasNext()) {
    2     input = scanner.nextLine();
    3     if (input.equals("0")) {
    4         break;
    5     }
    6     ...
    7 }

```

Fig. 4. Examples of RQ1 Java cases: a) Unused imports; b) Equals avoid null. The code formatting is unchanged.

A. RQ1: Most Prevalent Issues

Method. For each code quality issue with high frequency (from Figure 2), we analyzed its distribution among different tasks. We assume that if the distribution is biased towards a specific task, there should be some problem with the task itself. By examining the reason for this accumulation of issues in a particular task with concrete examples from the dataset, we can identify the true reason why students make this issue so often and whether it is their own fault or not. As examples, we selected and analyzed around 20 submissions from a concrete task that contains such an issue and the same amount where students avoid making it.

Case studies. In total, we highlight six cases: three for Java (Figure 4) and three for Python (Figure 5). Each case demonstrates how different factors, e.g., incorrect theory part, irrelevant tests, or complexity can influence code quality.

1) *Unused local variable (Top-1 Java issue):* This issue appears in 80% of the student submissions for Task#9057 [36], which asks to read several values, put them into placeholders in the text, and print the result. However, not all input values are required to complete the text, so if the student uses a separate variable for each of them, some remain unused. Some students avoid this issue by reading the value without saving it into a variable, but apparently, most of the students do not come to this idea because of the lack of experience in programming. We reported this observation to the creators of the task, now all the variables are used, and students no longer have such an issue in their submissions.

Summary. This example shows that the massive prevalence of this code quality issue was caused by poorly designed and tested tasks combined with the lack of knowledge by the students. Such kind of “third-party” mistakes, not related to the main idea of the task, can discourage the student and interfere with their learning of new complex material. Therefore, one possible solution to avoid this problem in this particular task would be to display this inspection without lowering the score, as optional or additional information, or even disable it.

2) *Unused imports (Top-2 Java issue):* This issue is common for all tasks, and most of it happens by accident because of the students’ carelessness. However, one example we would like to highlight is Task#3828 [37], where students need to find the maximum value in an array of numbers (Figure 4a). The task assignment explicitly says: “Try not to use the loop, but use Stream API”. Moreover, the template for the task contains

```

a) 1 print('
2 Did that stop the old Grinch?
3 No! The Grinch simply said,
4 "If I can't find a reindeer,
5 I'll make one instead!"
6 ')
Wrong multiline string usage

b) ...
for i in range(n):
    number = int(input())
    numbers.append(number)
Loop control variable not used

c) 1 print(1 * 2 * 3) ← Meaningless number operation violation

```

Fig. 5. Examples of RQ1 Python cases: a) Wrong multiline string use; b) Loop control variable not used; c) Meaningless number operation violation. The code formatting is unchanged.

import of `java.util.Arrays` to work with the Stream API. However, some students ignore this and write the code using a `for` loop, probably because they are more familiar with this construct. Conversely, more advanced students use `java.util.stream.Stream` library as an alternative for one mentioned above.

Summary. This is an example of how a code quality issue can indicate different intentions of teachers and students. The authors of the tasks expected only one particular solution from the students and configured the template only for that solution. However, they did not consider that some tasks can have alternative solutions, in which some parts of the template may not be needed, and therefore, cause an issue. In order to better control the submissions, the test base for these tasks should be strengthened or even extended with semantic checks.

3) **Equals avoid null:** This issue is interesting because it dominates in tasks with *moderate* complexity, although the peaks should logically occur in *shallow* tasks due to the inexperience of students. The cause of this peak is a specific *moderate* task, where 40% of all submissions with such an issue are concentrated. It requires the student to use the `equals` method to compare the input string with the constant string "0". However, the input value in Java can be `null`, so to avoid a `NullPointerException`, the best practice is to invoke `equals` on the constant, passing the input as an argument, rather than vice versa (see Figure 4b).

Summary. This task contains a theoretical section, the theory only provides examples of comparing two variables, so the student can find out about the practice of variable-to-constant comparison only from the Hyperstyle's feedback.

4) **Wrong multi-line string use (Top-3 Python issue):** The prevalence of this issue is mostly caused by Task#6881 [38], which requires students to print a multi-line string. In the first attempt, 88% of the students sent the straightforward submission with the multi-line string directly inside the `print` function (Figure 5a), which is flagged by Hyperstyle. Only a few students managed to avoid this issue by submitting the following solutions: (1) divide the string into several `print` calls; (2) use a single-line string with `\n` newline separators and invoke `print` only once; (3) put the text into a variable.

Summary. This task also contains a theoretical section. The theory part has a multi-line string usage example that contradicts the official Python style guide. When students try to use the incorrect code from the example, the issue is flagged.

5) **Loop control variable not used (Top-5 Python issue):** Task#6818 [39] requires students to write a program that

calculates the arithmetic mean. To solve the task, students add a loop without using its counter. In these cases, the variable should be replaced by an `_` symbol (see Figure 5b).

Summary. The theoretical section of this task does not have an example of this language feature. As a result, 48% of student submissions contain the issue in the first attempt. However, thanks to Hyperstyle, the students learn about it and in more than 90% of the cases get rid of it during their following attempts.

6) **Meaningless number operation:** This is another example of an issue where the peak in Figure 2 can be seen for the *moderate* tasks, which can be once again explained by specific tasks with moderate complexity causing it. Task#6558 [40] requires students to write a program that prints the product of these three numbers: `1 * 2 * 3` (see Figure 5c). However, according to the linter's rules, it is meaningless and could be simplified into `print(6)`. As a result, about 85% of these issues related to this particular task.

Summary. Obviously, the authors tried to artificially simplify the task for beginner students, missing the possible code quality issues. For this task, it would be better to use values from the input instead of the predefined constants, or put the code for obtaining input inside the pre-written template. This could probably keep the student's task simple but eliminate code quality issues. An alternative solution here could be to hide this issue from the student by customizing Hyperstyle for this particular task.

Discovered causes of code quality issues. (1) The task assignment is poorly developed and tested; (2) Students do not use constructs that are proposed by the creators; (3) The theoretical part of the task is incomplete or even incorrect; (4) Students have a lack of knowledge to avoid the issue.

B. RQ2: Dynamics of the Solution Quality

Method. To explain the differences in the dynamics of fixing issues, we also conducted a manual review of submissions, and collected some illustrative examples. For each issue, we were interested in the submission series of three different configurations: (1) only the first attempt contains an issue (the issue was fixed), (2) both the first and the last attempts contain an issue (the issue was not corrected), and (3) only the last attempt contains an issue (the issue appeared in the subsequent solutions). For this analysis, we selected tasks with the highest frequency of issues and analyzed their submission series.

Case studies. We considered four different cases with examples and explanations for several interesting trends, which we observed in Section IV. Figure 6 shows examples for the Java cases and Figure 7 — for the Python cases.

1) **Assignment in operand (Top-1 Java unfixed issue):** This issue often occurs in tasks where students are required to read numbers until a specific number is seen, e.g., Task#2153 [41]. There are several ways to implement this, but the most popular among students is to read the first number, assign it to the variable, and then read the remaining numbers into the same variable in the loop, until it equals the necessary value (see

```

a) 1 // Incorrect
   2 while ((x = getX() != 5)
   3     {...})
   4
   1 // Correct
   2 x = getX();
   3 while (x != 5) {
   4     ...
   5     x = getX();
   6 }

b) 1 // Incorrect
   2 if ((x = getX() == 5)
   3     {...})
   4
   1 // Correct
   2 x = getX();
   3 if (x == 5) {...}

c) 1 // Incorrect
   2 ExecutorService executor
   3     = Executors
   4     .newCachedThreadPool();
   5
   1 // Correct
   2 ExecutorService executor
   3     = Executors
   4     .newCachedThreadPool();

d) 1 // Incorrect
   2 for (...) {...};
   3
   1 // Correct
   2 for (...) {...}

```

Fig. 6. Examples of RQ2 Java cases: a, b) Assignment in operand; c, d) Empty statement check.

Figure 6a). To get rid of code duplication, experienced students start to save the numbers into a variable directly inside the `while` condition (Figure 6a and Figure 6b), which leads to an *Assignment in operand* issue. It is worth noting that this code is also correct and is even sometimes used for size optimization of the bytecode. However, in industrial programming, this code construction is considered to be error-prone, as it is difficult to read and understand.

Summary. The theory on the platform does not provide such a complex example of the `while` usage. However, we found many students that have only two submissions with a small difference in time (within 5 minutes) and the second one has this complicated case with the `while` usage. Taking all facts into account, we can assume that these students, after successfully submitting a solution, used the opportunity to see other students' successful public solutions on the JetBrains Academy platform, saw this interesting new language construct with the `while` usage, copied it into their code and submitted it. This example shows how a seemingly useful platform feature can have an unexpected effect in the form of spreading a mistake.

2) *Empty statement check (Top-2 Java fixed issue):* We found three typical instances of this issue. Firstly, there are two semicolons one after another next to a variable initialization. The task has a pre-written template with a variable declaration with a semicolon at the end. When students initialize the variable with a value, they might also add their own semicolon automatically, forgetting about the one already present in the template (see Figure 6c).

Secondly, there can be a redundant semicolon after a control flow statement. Students are taught that in Java it is necessary to put a semicolon after each statement, so diligent students also do this after the body of `if`, `for`, and `while` blocks (Figure 6d), although it is not required.

Finally, there can be a control flow statement without a body like `for (i = 2; i <= n; i *= 2);`. Often-times, students try to shorten their code and calculate some value directly using a `for` loop variable, leaving the body of the `for` statement empty. For example, this can happen to calculate the power of some value. However, this is an anti-

```

a) 1 # Incorrect
   2 ...
   3 t = r"..."
   4 match = re.match(t, string)
   5 ...
   6 print(f"Full number: "
   7       f"{match.group(1)}")

b) 1 # Incorrect
   2 print(f'Amount'
   3       f'={int(round(tax_amount))}')
   4
   1 # Correct
   2 tax_amount = int(round(tax_amount))
   3 print(f'Amount={tax_amount}!')

c) 1 # Incorrect
   2 print(a < b and b < c)
   3
   1 # Correct
   2 print(a < b < c)

```

Fig. 7. Examples of RQ2 Python cases: a, b) Too complex formatted string violation; c) Chained comparison.

pattern, since an empty `for` body often introduces bugs that are hard to spot later on.

Summary. Most students (over 85%) fix this issue because it seems pretty easy to fix — simply remove the extra semicolon. However, about 15% of users do not to fix this issue. A possible reason for this may be that the issue is too simple and students prefer to fix more interesting ones.

3) *Too complex formatted string (Top-1 Python unfixed issue):* This issue appears if the student uses an `f-string` and calls something inside it. For instance, the example in Figure 7a calls the `group` function, and the example in Figure 7b calls the `int` and `round` functions. This can complicate the code and lead to potential bugs.

Summary. Almost all students write the correct solution at the beginning, but then, in an attempt to shorten it, they transfer some of their code into the `f-string`, which causes the issue. At this stage, students might not yet be familiar with the concept of refactoring and therefore do it incorrectly.

4) *Chained comparison (Top-3 Python fixed issue):* This issue appears if students use two comparisons instead of a single triple one (see Figure 7c). Python allows chaining comparison operators as a helpful shorthand, but such a “syntactic sugar” is not common in other programming languages.

Summary. There are not many guidelines about code quality on the JetBrains Academy platform at the moment. However, Chained comparison issue is covered by the theory in Task#5920 [42]. This task provides several illustrative examples and clear explanations, which can be the reason why 83% of the students fix it in their final attempts.

Discovered reasons to fix code quality issues. (1) When students exchange their solutions, they can both acquire new knowledge and borrow other students' issues; (2) If an issue is simple to correct, students mostly try to do it, even with issues that are not covered in the theory part; (3) However, if the issue is too difficult to understand, students may ignore it. In this case, theory can help.

VI. ANOMALOUSLY LONG SUBMISSION SERIES

In the main body of this study, answering two research questions above, we did not take into account long submission series (longer than 5 successful submissions), which constitute about 5% of the entire dataset. In this additional study, we



Fig. 8. The full submission series for solving Task#9261 [43], consisting of 39 attempts. The green color indicates places with the fixed code quality issues between attempts. Comments contain the list of code quality issues in the code line and their number in the brackets (if it contains more than one such issue). (1) Contains the first successful attempt, (4) — the last attempt. Between (2) and (3) there were 35 attempts, the changes are shown next to the corresponding arrow, changing line 10.

manually analyze these long sequences to examine how students fix issues and why they might make numerous attempts. In the previous sections, we omitted code quality issues related to formatting, because they are less interesting and they would dominate the results in RQ1. However, in this part, we kept all formatting issues because they play a major role in these anomalously long sequences.

A. Data

For this analysis, we extracted all student submission series of length greater than 5 from the dataset described in Section III. Then, we filtered out submissions made by internal JetBrains Academy bots, which resend all submissions for a specific task after some content or any of the tests change. From the standpoint of out dataset, these bots look like separate students that have suspiciously long artificial submission series of more than 1,000 attempts, which were sent one by one with less than a minute gap, but had quite different code inside (solutions of different students), making them easy to detect and remove. Thus, the final data for this analysis consisted of only 5,180 submission series in Java and 2,786 in Python, with 5 to 50 attempts in each.

Finally, we calculated the frequency of different lengths of the remaining series. For Java, most students have 6-10 attempts (about 92.6%). However, several students solve tasks in 31, 39, and even 41 attempts. For Python, the situation is similar — about 93.9% of the students have 6–10 attempts.

However, the maximum number of attempts for Python is slightly higher — one of the students tried to pass the solution 47 times. At the same time, the rest of the students made no more than 29 attempts, which is lower than in Java.

B. Java Case Study

We chose the more interesting of the longest series to present as a case study. This submission series has 39 attempts for Task#9261 [43], where it is necessary to read several lines from the keyboard and output them in reverse order. The full history of correct submissions is presented in Figure 8.

The first successful attempt that passes all tests (1) has 6 unique code quality issues and 19 issues in total. As we can see, almost all of them are about code formatting in `for` loops. To begin with, the student got rid of indentation issues step by step within 3 attempts (2): (1) firstly, deleted extra space between the open bracket and `int` in lines 10 and 18; (2) then added missing spaces in the `guestsArr` array index in line 19; (3) finally, removed extra spaces between `i++`, `j++`, `i--`, and the closing bracket in lines 10, 12, and 18. We assume that the student knew exactly how to fix the issue at each attempt, as it was fixed in several places via one change.

However, after these changes, the `for` loops still contains formatting issues: *Whitespace After* and *Whitespace Around*. In fact, the loops do not contain a space after the `for` keyword and before the opening curly bracket. Probably, because the JetBrains Academy interface highlights only the line with the

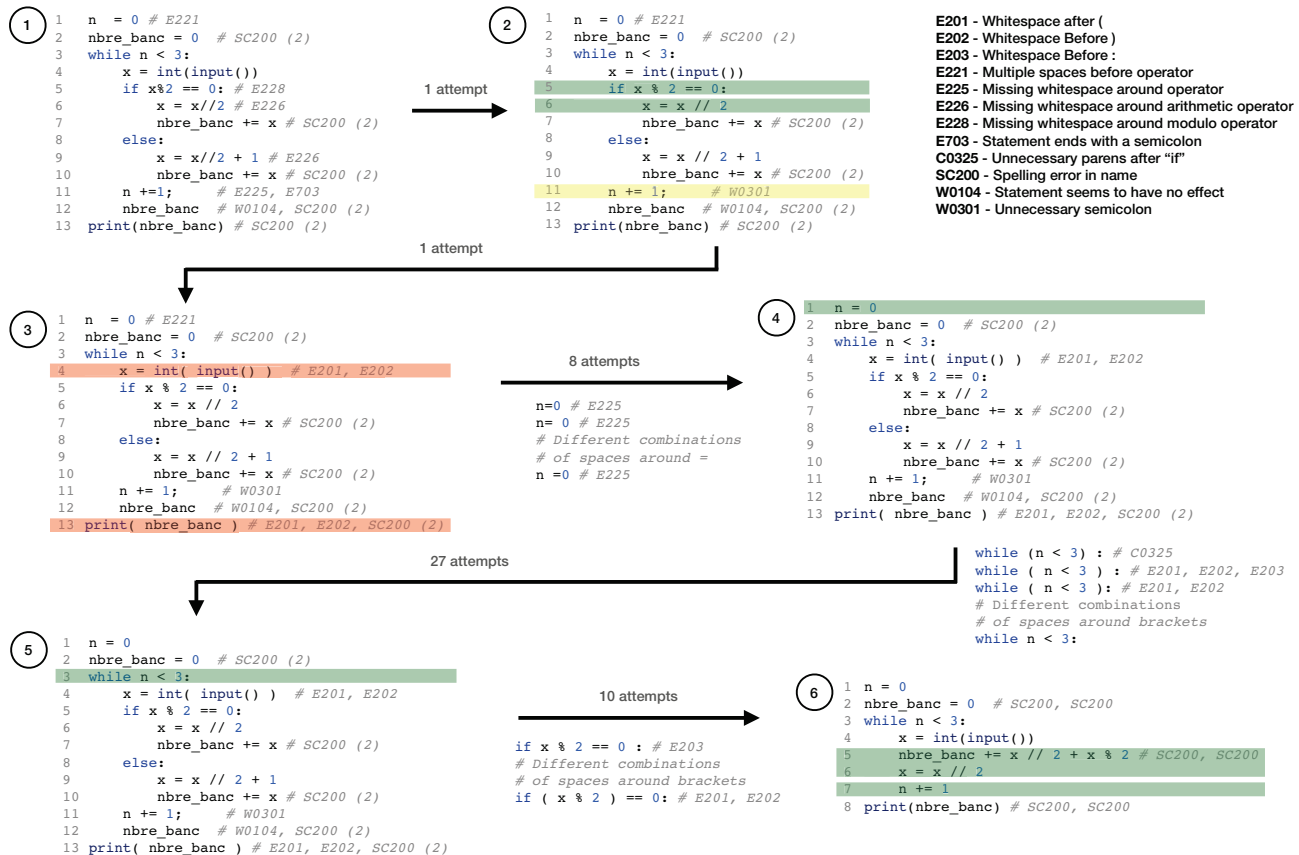


Fig. 9. The full submission series for solving Task#6462, consisting of 47 attempts. The green color indicates places with the fixed code quality issues between attempts. The yellow color indicates places where the initial issues were corrected, but new ones appeared. The red color indicates places where new code quality issues appeared. Comments contain the list of code quality issues in the code line and their number in the brackets (if it contains more than one such issue). (1) contains the first successful attempt, (6) — the last attempt. Between (3) and (4) there were 8 attempts, the changes are shown next to the corresponding arrow, changing line 1. Between (4) and (5) there were 27 attempts, the changes are shown next to the corresponding arrow, changing line 3. Between (5) and (6) there were 10 attempts, the changes are shown next to the corresponding arrow, changing line 5.

issue, the student did not understand their exact positions and went through various options for 35 attempts. Finally, the student found the right combination (3). In the end, a more serious issue was fixed — the `for` loop was replaced with `forEach` (4). The final code looks more readable, but still contains a few issues — all of them do not affect the final grade, so the student finished this task with the highest score.

C. Python Case Study

For Python, we analyzed the longest submission series with 47 attempts. The student tried to solve Task#6462 [44], which requires reading three integer numbers from the user input (the number of students in three classes) and calculating the minimum number of desks to be purchased if at most two students may sit at any desk. The full history of the correct submissions is presented in Figure 9.

The first successful attempt that passes all the tests (1) has 7 unique code quality issues and 17 issues in total, and almost all of them are about code formatting in arithmetic expressions. Thus, the student tried to correct them first and succeeded (lines 5, 6 and 11 in (2)). However, another issue in line 11 (E703) was unexpectedly substituted with (W0301).

Both of these issues point to an extra semicolon at the end of line 11, but they were detected by different linters inside the Hyperstyle tool [27] and deduplicated in a wrong way.

Probably, this platform behavior confused the student, so in the third attempt (3), the student added extra spaces around `input` (line 4) and inside `print` (line 13), which produced four new formatting issues. Then, the student switched to the first line and tried to correct wrong spaces in 8 attempts (4). It is interesting to note that the student tried different combinations of spaces, while there were already similar lines in the code that did not contain this issue (e.g., lines 2, 4, 6, etc.). Besides the highlighting of the exact issue position, a clearer message from the tool could help with this task faster, since the current version of the hint is too general: *missing whitespace around operator*.

Next, the student added parentheses around the `while` statement and tried to add spaces in the correct way. After 27 attempts, the extra brackets were removed and the solution between (4) and (5) did not change. Finally, we see the same problem, but with the `if` statement between (5) and (6). For both of these cases, the student tried to correct the issues that

they themselves added, experimenting with parentheses and spaces. Perhaps, the problem is that students can not see their submissions history and roll back code with code quality issues they made to the previous state.

As a result, after all corrections, the student's final solution (6) is much cleaner and shorter than in (1). However, it should be noted that issues related to incorrect spelling and those that do not affect the final score have not been corrected.

Summary. (1) Issue hints must be carefully adapted in code quality assessment tools because students can spend a lot of attempts to fix an issue just by going through all the possible options; (2) Students can change correct code to incorrect code. Perhaps, in this case, it is necessary to additionally indicate these cases to the student and highlight them among other issues; (3) Some students find it difficult to correct formatting issues (spaces, brackets, etc) in the Web editor if the exact position is not set, resulting in a lot of attempts; (4) Students are reluctant to correct issues that do not affect the final code quality grade.

VII. IMPLICATIONS

Let us reiterate the main practical results from the analysis of cases in both RQs and anomalously long submission series.

Theoretical part. The theoretical part of a task should be complete and covering different cases. First and foremost, it should not contain quality issues in itself, and, equally important, it should cover all the specific sub-cases that will appear in the practical part.

Extensive testing. Test cases for practical tasks should be exhaustive and account for different possible solutions. Perhaps, this is not feasible to predict from the start, so this part should be revisited after some time by analyzing the most popular solutions.

Pre-written templates. If the task contains a pre-written template, the template also should not contain code quality issues itself. Moreover, the template should not be too restrictive and force a particular solution, but instead enable different solution strategies.

Reporting issues. The overall results of our study indicate that the reporting of code quality issues to students is paramount in the education process. In a lot of cases, these reports allow the students to fix the issues in their initial attempts, thus practically instilling code quality guidelines.

Rewarding students. Specifically, we found that the *grade* for code quality (that does not directly influence the overall result of the course) is a good incentive for a student to put effort into the task, and, conversely, some issues that do not influence the grade, remain unfixed.

Adapting the hints. As mentioned before in Section II, a problem with some code quality tools is that their results are unadapted for novices. In our study, we also found that students sometimes struggle with understanding what exactly the issue is. Overall, even when adopting messages from tools, developers should be very careful with wording. Also, for simple and understandable formatting issues, finding their exact place in the marked line can also present a challenge.

VIII. THREATS TO VALIDITY

In this work, we conducted a large-scale study of code quality issues, however, it has some limitations. In RQ1, the presence of each specific issue is not uniform among different tasks, which means that there are tasks where a specific issue is more prevalent, and these tasks can inflate its prevalence in general. However, this is normal: different tasks target different constructs and concepts in the language, and are thus more prone to different types of issues. For this reason, issues such as *Missing break in switch* are more prevalent in certain tasks, while more general issues like *Redundant import check* are more universal. We believe that this does not invalidate the results of our study, since we used a diverse dataset of 779 tasks that target a wide variety of topics.

Also, the results of this work are not generalizable to all MOOC platforms, since they were obtained on a specific platform — JetBrains Academy. We conducted the manual analysis to look into case studies for RQ1 and RQ2, as well as long series, however, these findings are not directly generalizable, since they depend on specific tasks and environments. At the same time, we received and shared a lot of insights that can be useful to the community in different settings.

IX. CONCLUSION AND FUTURE WORK

In this paper, we carried out an analysis of the code quality of successful student submissions from the JetBrains Academy platform. We studied the most popular code quality issues among the users' initial successful submissions, analyzed the dynamics of fixing these issues on a large scale, and analyzed the most interesting cases manually.

We found that students introduce various kinds of code quality issues, with the real root causes being not only the students themselves, but also external factors, such as an incorrect task or learning materials, as well as obscure code quality issues messages. The issues also significantly differ in terms of how often they get fixed. Some issues are fixed by the majority of students if the issue is simple to understand or easy to fix. However, some issues are fixed more rarely, and there even exist issues that are more common in the last attempts than in the first. This can happen due to students copying other students' solutions after solving the task, or when they try to shorten the code without the proper knowledge of refactoring practices. Finally, in about 5% of the cases, the submission series can be more than 5 attempts long, and we found that these long series of attempts often consist of unsuccessfully fixing simple formatting issues. The supplementary materials for the paper are available online [28].

We believe that this study will be useful to teachers, developers of code quality tools, as well as MOOC educational platforms in general. One can focus on specific issues highlighted in this work to make sure they are addressed, but it is also crucial to improve existing materials and tools, in particular so that they do not introduce issues themselves. In future work, we plan to study in greater detail the influence of grading on code quality, as well as differences in code quality issues between the Web environment and the IDE.

REFERENCES

- [1] B. Boehm, "A view of 20th and 21st century software engineering," in *Proceedings of the 28th international conference on Software engineering*, 2006, pp. 12–29.
- [2] H. Zhong and Z. Su, "An empirical study on real bug fixes," in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, vol. 1. IEEE, 2015, pp. 913–923.
- [3] J.-L. Lions, L. Luebeck, J.-L. Fauquembergue, G. Kahn, W. Kubbat, S. Levedag, L. Mazzini, D. Merle, and C. O'Halloran, "Ariane 5 flight 501 failure report by the inquiry board," 1996.
- [4] R. C. Martin, *Clean code: a handbook of agile software craftsmanship*. Pearson Education, 2009.
- [5] R. L. Glass, *Facts and Fallacies of Software Engineering*. Addison-Wesley Professional, 2002.
- [6] M. Fowler, *Refactoring: improving the design of existing code*. Addison-Wesley Professional, 2018.
- [7] M. Fagan, "Design and code inspections to reduce errors in program development," in *Software pioneers*. Springer, 2002, pp. 575–607.
- [8] H. Keuning, B. Heeren, and J. Jeuring, "Code quality issues in student programs," in *Proceedings of the 2017 ACM Conference on Innovation and Technology in Computer Science Education*, 2017, pp. 110–115.
- [9] J. Börstler, H. Störrle, D. Toll, J. Van Assema, R. Duran, S. Hooshangi, J. Jeuring, H. Keuning, C. Kleiner, and B. MacKellar, "'I know it when I see it' Perceptions of Code Quality: ITiCSE'17 Working Group Report," in *Proceedings of the 2017 ITiCSE Conference on Working Group Reports*, 2018, pp. 70–85.
- [10] H. Keuning, B. Heeren, and J. Jeuring, "How teachers would help students to improve their code," in *Proceedings of the 2019 ACM Conference on Innovation and Technology in Computer Science Education*, 2019, pp. 119–125.
- [11] E. G. Oh, Y. Chang, and S. W. Park, "Design review of MOOCs: Application of e-learning design principles," *Journal of Computing in Higher Education*, vol. 32, no. 3, pp. 455–475, 2020.
- [12] C. Impey and M. Formanek, "MOOCs and 100 Days of COVID: Enrollment surges in massive open online astronomy classes during the coronavirus pandemic," *Social Sciences & Humanities Open*, p. 100177, 2021.
- [13] A. Koutropoulos, M. S. Gallagher, S. C. Abajian, I. de Waard, R. J. Hogue, N. O. Keskin, and C. O. Rodriguez, "Emotive Vocabulary in MOOCs: Context & Participant Retention," *European Journal of Open, Distance and E-Learning*, 2012.
- [14] S. Kinash, "MOOCing about MOOCs," *Education Technology Solutions*, vol. 57, no. 70, pp. 56–58, 2013.
- [15] J. Carter, K. Ala-Mutka, U. Fuller, M. Dick, J. English, W. Fone, and J. Sheard, "How shall we assess this?" in *Working group reports from ITiCSE on Innovation and technology in computer science education*, 2003, pp. 107–123.
- [16] G. De Ruvo, E. Tempero, A. Luxton-Reilly, G. B. Rowe, and N. Giacaman, "Understanding semantic style by analysing student code," in *Proceedings of the 20th Australasian Computing Education Conference*, 2018, pp. 73–82.
- [17] S. H. Edwards, N. Kandru, and M. B. Rajagopal, "Investigating static analysis errors in student Java programs," in *Proceedings of the 2017 ACM Conference on International Computing Education Research*, 2017, pp. 65–73.
- [18] I. Albluwi and J. Salter, "Using static analysis tools for analyzing student behavior in an introductory programming course," *Jordanian J. Comput. Inf. Technol*, 2020.
- [19] E. Aivaloglou and F. Hermans, "How kids code and how we know: An exploratory study on the Scratch repository," in *Proceedings of the 2016 ACM Conference on International Computing Education Research*, 2016, pp. 53–61.
- [20] P. Techapalokul and E. Tilevich, "Understanding recurring quality problems and their impact on code sharing in block-based software," in *2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 2017, pp. 43–51.
- [21] Y. Bai, T. Wang, and H. Wang, "Amelioration of teaching strategies by exploring code quality and submission behavior," *IEEE Access*, vol. 7, pp. 152 744–152 754, 2019.
- [22] N. C. C. Brown, M. Kölling, D. McCall, and I. Utting, "Blackbox: A large scale repository of novice programmers' activity," in *Proceedings of the 45th ACM technical symposium on Computer science education*, 2014, pp. 223–228.
- [23] "Scratch Open Dataset," <https://scratch.mit.edu/explore/projects/all/>, [Online; accessed 20-January-2023].
- [24] C.-A. Lo, Y.-T. Lin, and C.-C. Wu, "Which programming language should students learn first? A comparison of Java and Python," in *2015 International Conference on Learning and Teaching in Computing and Engineering*. IEEE, 2015, pp. 225–226.
- [25] D. Liu and A. Petersen, "Static analyses in Python programming courses," in *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*, 2019, pp. 666–671.
- [26] "JetBrains Academy," <https://www.jetbrains.com/academy/>, [Online; accessed 20-January-2023].
- [27] A. Birillo, I. Vlasov, A. Burylov, V. Selishchev, A. Goncharov, E. Tikhomirova, N. Vyahhi, and T. Bryksin, "Hyperstyle: A tool for assessing the code quality of solutions to programming assignments," in *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education V. 1*, 2022, pp. 307–313.
- [28] "Artifacts and supplementary materials," <https://zenodo.org/record/7573422>, [Online; accessed 20-January-2023].
- [29] T. Effenberger and R. Pelánek, "Code quality defects across introductory programming topics," in *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education V. 1*, 2022, pp. 941–947.
- [30] A.-J. Molnar, S. Motogna, and C. Vlad, "Using static analysis tools to assist student project evaluation," in *Proceedings of the 2nd ACM SIGSOFT International Workshop on Education through Advanced Software Engineering and Artificial Intelligence*, 2020, pp. 7–12.
- [31] H. Keuning, B. Heeren, and J. Jeuring, "Student refactoring behaviour in a programming tutor," in *Koli Calling'20: Proceedings of the 20th Koli Calling International Conference on Computing Education Research*, 2020, pp. 1–10.
- [32] "PMD," <https://pmd.github.io/>, [Online; accessed 20-January-2023].
- [33] S. H. Edwards and M. A. Perez-Quinones, "Web-cat: automatically grading programming assignments," in *Proceedings of the 13th annual conference on Innovation and technology in computer science education*, 2008, pp. 328–328.
- [34] "Pylint," <https://pylint.pycqa.org/en/latest/>, [Online; accessed 20-January-2023].
- [35] "JetBrains," <https://www.jetbrains.com/>, [Online; accessed 20-January-2023].
- [36] "Task#9057 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/9057>, [Online; accessed 20-January-2023].
- [37] "Task#3828 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/3828>, [Online; accessed 20-January-2023].
- [38] "Task#6881 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/6881>, [Online; accessed 20-January-2023].
- [39] "Task#6818 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/6818>, [Online; accessed 20-January-2023].
- [40] "Task#6558 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/6558>, [Online; accessed 20-January-2023].
- [41] "Task#2153 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/2153>, [Online; accessed 20-January-2023].
- [42] "Task#5920 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/5920>, [Online; accessed 20-January-2023].
- [43] "Task#9261 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/9261>, [Online; accessed 20-January-2023].
- [44] "Task#6462 on the JetBrains Academy platform," <https://hyperskill.org/learn/step/6462>, [Online; accessed 20-January-2023].