



Automated Questions About Learners' Own Code Help to Detect Fragile Prerequisite Knowledge

Teemu Lehtinen
teemu.t.lehtinen@aalto.fi
Aalto University
Espoo, Finland

Otto Seppälä
otto.seppala@aalto.fi
Aalto University
Espoo, Finland

Ari Korhonen
ari.korhonen@aalto.fi
Aalto University
Espoo, Finland

ABSTRACT

Students are able to produce correctly functioning program code even though they have a fragile understanding of how it actually works. Questions derived automatically from individual exercise submissions (QLC) can probe if and how well the students understand the structure and logic of the code they just created. Prior research studied this approach in the context of the first programming course. We replicate the study on a follow-up programming course for engineering students which contains a recap of general concepts in CS1. The task was the classic rainfall problem which was solved by 90% of the students. The QLCs generated from each passing submission were kept intentionally simple, yet 27% of the students failed in at least one of them. Students who struggled with questions about their own program logic had a lower median for overall course points than students who answered correctly.

CCS CONCEPTS

- **Applied computing** → **Interactive learning environments;**
- **Social and professional topics** → **Computer science education.**

KEYWORDS

QLC; prerequisite knowledge; program comprehension; online education

ACM Reference Format:

Teemu Lehtinen, Otto Seppälä, and Ari Korhonen. 2023. Automated Questions About Learners' Own Code Help to Detect Fragile Prerequisite Knowledge. In *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1 (ITiCSE 2023)*, July 8–12, 2023, Turku, Finland. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3587102.3588787>

1 INTRODUCTION

While novice programmers might be able to produce a working solution to a programming problem, some of them might struggle and answer incorrectly even simple questions about their own (functionally correct) program code [8]. Previous research suggests that this phenomena may be due to fragile programming skills in general [13], misconceptions on programming constructs [11, 20], inability to trace code [1, 10] or that the failing students did not construct the program by themselves [22]. Recent advances in AI

are likely to exacerbate the problem by supplying code that the student may not fully comprehend. All of the previous reasons can seriously hinder learning on a second programming course.

We partially replicate previous studies on automatically generated *Questions about Learners' Code* (QLCs) [7, 8, 16] by adapting selected research questions to new context. The previous studies focused on students' program comprehension during the first weeks of their *first programming course* (CS1). The studies found that a number of students perform poorly when answering questions on their own functionally correct code. In addition, two studies reported that students who answer incorrectly to QLCs have on average lower success on the course.

The three papers call for replications and improved research designs as future efforts. In the present study we answer this call: research QLCs in a new context, a programming course for non-CS majors with a CS1 as prerequisite. As the students are further in their programming studies we expect them to have better programming skills and ability to answer QLCs similar to those in the previous studies. We do not, however, have access to their grades or research data about their performance on the prerequisite course.

We utilize QLCs on a *second programming course* to collect cases where students submit a program, but do not properly understand how it works. A variation of the classic Rainfall-problem [21] was used as a programming recap exercise. The students submitted their solutions to an automated grading platform allowing resubmits. After the programming task, students were offered multiple-choice QLCs that were generated for the program they submitted and which we expected the programmer should be able to answer. The students' answers to these questions were analyzed to discuss the following research questions:

RQ1 How well do students answer QLCs for an exercise targeting prerequisite knowledge?

RQ2 How do the QLC results relate to success on second programming course?

We hypothesized that QLCs could reveal weaknesses in prerequisites which again can lead to low grades or even drop-outs. We found that success rates for the different types of QLCs were comparable to those reported in previous studies [7, 8, 16] even though the students had passed a CS1 in Python. Furthermore, students answering incorrectly about their own program logic had a lower median for course points compared with those that answered correctly.

The rest of the paper is structured as follows. Section 2 gives some background and introduces the related studies of which selected questions are replicated in this research. Section 3 gives the context in which this study is conducted, introduces the exercise used in the study, and explains how the data was analysed. Section 4



This work is licensed under a Creative Commons Attribution International 4.0 License.

ITiCSE 2023, July 8–12, 2023, Turku, Finland
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0138-2/23/07.
<https://doi.org/10.1145/3587102.3588787>

summarises the results, and finally Section 5 discusses the findings, their interpretations, and some threats to validity.

2 RELATED WORK

Program comprehension (PC) has been researched extensively. Schulte et al. [18] review PC literature from a computer science education research perspective. They map existing PC models to a proposed *revised block model* for education. The block model describes PC in two dimensions: a scale of atoms–blocks–relations–macro and other dimension on text–execution–purpose [17]. Those dimensions combine to areas or a diagram of blocks where we can examine comprehension; for example, *execution* through a sequence of *related* function calls.

Although many people often think of programming with focus on writing code, researchers have recognized the importance of teaching program comprehension to novices and many learning activities are suggested and developed to address this issue [3, 4, 6, 12, 19]. Another common conception is that if a person writes a program they can reason how it works. However, this is not always true and cases where students have poor comprehension of the programs they wrote have been reported [5, 8, 11].

Perkins and Martin [13] researched difficulties that novice programmers encounter. In their experiment, instructors worked one-on-one with the programming students and asked them questions such as what does this element do? While such prompts help students go forward the students' reactions also reveal cases of fragile knowledge. They further report on different types of fragility they identified.

Lehtinen et al. [9] defined QLCs as an approach that poses questions to a student about the concrete structures and patterns in a program, which the same student previously created. They hypothesize such questions could test different areas of PC, catch cases of unproductive success, and trigger reflection including self-explanation. Furthermore, they describe a potential design of a system to automate generation of QLCs.

Lehtinen et al. [8] report on an early pilot for QLCs where students wrote open-text answers to QLCs that were manually prepared for selected program writing exercises for novices. Their main research questions are “How well do students answer QLCs?” and “How success in QLCs correlates with other learning data?”. They used the block model to design their QLCs and report the following student success rates: atom–text 95%, relation–text 67%, atom–execution 72%, block–execution (80% and 75%), macro–execution (58% and 79%), and macro–purpose 91%. Furthermore, among the students who created a functionally correct program the ones who answered the related set of the first three QLCs correctly had increased course success and retention in comparison to students who answered incorrectly.

Santos et al. [16] present a tool that automates QLCs for Java programs and report on a pilot evaluation on CS1. The evaluation asks “How do students perform in QLCs in their lab exercises?” and “How do students perceive the activity of answering QLCs?”. They measure above 80% success rates for QLCs targeting static aspects (text in block model) and below 50% success rates for QLCs targeting dynamic aspects (execution in the block model). Despite

the low success in the latter, students reported good average confidence levels in both static and dynamic cases. The authors also note that success rates have large variations for identical QLC templates when the programming task changes, in other words, difficulty is dependent on the program. In a post-exercise survey majority of the students reported learning or reinforcement of their knowledge regarding terminology. Approximately one third of students reported the same regarding loops, recursion, and variable values.

Lehtinen et al. [7] present automated QLCs for JavaScript programs as well as pilot study on students' answers. They report success rates all the way from 33% up to 97%. In the upper end the QLCs target block model's text level and in the lower end they require tracing in block model's execution level. They find that students who repeatedly answer QLCs incorrectly also have more challenges while writing the program and tinker towards solutions. Respectively, the average error rate in QLCs correlates negatively with course success.

3 METHOD

3.1 Context

This study was conducted in a 12 weeks 5 ECTS¹ programming course given in Aalto University, a large research university in Finland. The course includes topics such as introduction to Numpy and Scipy modules in Python, basics in databases and how to interconnect Python and SQL, MATLAB and Simulink programming, and simple web programming with Javascript. The course has one 5 ECTS prerequisite programming course in Python. QLCs were designed to be part of a recap exercise in week 1, which also included three other programming exercises and a couple of multiple choice questions. All of the exercises were automatically assessed and the students got immediate feedback on their solutions. In case the student did not achieve full points on the first attempt, 9 resubmissions were allowed.

The course had no final exam, but the final grade was calculated based on the points received from the exercises and a project done in groups (max. 3 students). To pass the course, students needed to do at least half of the individual weekly exercises and the project. To support this, the course had lab sessions on every weekday for seeking help from teaching assistants (TA). The course had also 5 live lectures and a couple recorded video lectures. Due to the COVID-19 situation, the live lectures and lab sessions were in Zoom².

Typically some 350 engineering students³ start doing the course each year. There are no penalties for dropping the course and only about 75-80% pass the course yearly. Most of the students are first year students, but many might end up taking the course on the second or third year, which also explains the relatively high drop out rate.

The study was conducted in Spring 2021. 324 students volunteered to participate and gave their consent on the online course

¹The European Credit Transfer and Accumulation System

²Zoom is a communications platform that allows users to connect with video, audio, phone, and chat over Internet.

³Students majoring e.g. from Mechanical and Civil Engineering, but not Computer Science

platform. While all students did the same tasks, the volunteering students allowed their anonymous exercise data (timestamps, answers, automated feedback, points) to be used in our research. We offered no incentives for the participation and informed that participation had no effect on the course requirements or assessment. The course instructor was one of the second authors. The local regulations did not require an ethical board's review for our research.

3.2 Task

The first week of the course included Python recap-material. We selected the first program writing exercise, called "Calculation of average rainfall", for our QLCs experiment. For the rest of the course and the 98 other exercises we only calculate students' exercise points. The rainfall task was to create a program according to the following description which we translated to English for this article:

"Implement function rain(), which asks the user for days' rainfalls and calculates the average rainfall for the given days. The program ends by providing -999 as input. Notice that rainfall cannot be negative so your program must ignore negative numbers. The program must also ignore other inputs beside numbers (e.g. letters). If the user does not input any acceptable numbers the program must return 0."

In addition to the description, an example of expected input and output lines was given. The provided code template included an empty function stub and a "main" block.

The students submitted their program online and received automated feedback based on four unit test cases:

- (1) the program ends when -999 is given as input
- (2) the program does not fail when lines of letters is given as input,
- (3) the program ignores input that cannot be transformed into integer numbers,
- (4) the program outputs the average of the numbers given as input.

The differences from the expected output were displayed to the student. In addition, TAs were available to answer questions and give hints during online lab sessions. Based on the number of passed unit test cases students received up to 95 points from the program writing task. Students could resubmit their program up to 9 times to receive new assessment and points.

Once the students had submitted their program at least once they could open a questionnaire that included 2–3 generated QLCs. If the program did not include a structure necessary to create one of the question types the student would not receive that question type. Students could only answer the QLCs once and received 5 points if they answered correctly. Together the program writing and QLCs were worth 100 points which were 1.4% of the total available on the course.

QLC generation uses question templates and static program analysis as described in [9]. The library for generating QLCs for any python code is available as open source software⁴ and contains both documentation and code examples. The selection of QLC types has been further extended from the types included in this study.

Figure 1 presents an example questionnaire that was generated for the included program and then answered by the user. These were multiple-choice questions and students could only answer

* Which of the following are variable names in the program? **Correct**

- ☐ except A reserved word in the programming language
- ☒ sum A variable in the program
- ☐ float A function that is defined internally in the programming language
- ☐ important An unused word
- ☒ i A variable in the program

* From which line can program execution jump to line 12? **Incorrect**

- ☒ From line 6 At minimum the first line inside a try-block will start to execute
- ☐ From line 10 This line should not cause a compatible error
- ☐ From line 8 This line can cause a compatible error
- ☐ From line 5 Except-block is unreachable from outside the related try-block

* Which of the following best describes the purpose of line 14? **Incorrect**

- ☐ Is a condition for ending the program
- ☐ Accepts new data
- ☐ Guards against division by zero **Correct**
- ☒ Ignores negative input

Submit

```

1 def rain():
2     sum = 0
3     i = 0
4     row = 0
5     while row != -999:
6         try:
7             row = input("Enter rainfall as millimeters.
8             Stop by entering -999.\n")
9             row = float(row)
10            if row >= 0:
11                sum += row
12                i += 1
13            except ValueError:
14                continue
15        if i > 0:
16            avg = sum / i
17        else:
18            avg = 0
19    return avg

```

Figure 1: An example of a generated and answered questionnaire for a program that completes the rainfall task.

once. For increased learning, descriptions for each option were displayed after the user had submitted their answers.

Table 1 presents the selected QLC templates (Q1-Q3) and descriptions of their correct answer options as well as their incorrect distractor options. The QLCs were designed to approach different areas of program comprehension. We situate these questions in the two dimensions of the block model for program comprehension.

Q1 has been a typical question type in previous QLC research. It tests the students' recognition of an atomic element, in this case a variable, in the program's structure (atom-text).

Q2 asks from which lines can the execution enter an except-block. While testing understanding of block-structures, it relates the expected error to the potential source of the error (relations-execution).

Q3 takes a step deeper into the purpose of the code. More specifically, it asks for the purpose of a statement in its context (atom-purpose). The answer options are designed so that students should

⁴<https://github.com/teemulehtinen/qlcqp>

have some understanding of the program’s design to answer correctly.

3.3 Analysis

Anonymized data was extracted from the course platform for those 324 students that gave a research consent. We rejected 33 students who did not create the rainfall program or did not answer the related QLCs. Finally, we accepted 291 students (90%) in our research and analysed their exercise data quantitatively.

To answer RQ1, we counted the numbers of correct as well as incorrect answer options. For each type of question, we calculated success rates as the number of correct answers divided by the total number of answers. We discuss the relative difficulty of the different QLC types in this study and compare to success rates reported in earlier research.

To measure success on course we use *Course Points* that is the sum of exercise points student earned during the course. The course grade was assigned based on these points as well. To answer RQ2, we separately investigate the three types of QLCs we generated. For each QLC type, we compare the students who answered correctly to those who answered incorrectly.

According to visual inspection, the course points have non-normal distribution. The maximum points available on the course cause a ceiling effect as the automatically assessed exercises can be submitted 10 times to get feedback and collect more points. Also course points tend to accumulate close to the limits required for different course grades.

In case of non-normal distributions, we calculate Mann-Whitney *U*-tests for the hypothesis that two groups, students answering QLCs correctly and incorrectly, have equal medians. In addition, we calculate common language effect size. We use an alpha level of $p < .05$ to reject the hypothesis and as we test three comparisons we use Bonferroni correction $p/3 = .017$ to avoid false-positive analysis.

4 RESULTS

Out of the 291 students in the study 79 students (27%) failed at least one of the QLCs generated for their program. Table 1 presents success rates and counts of incorrect answers that students have for each of the three different QLC types. Next, we highlight the findings for each type and include the question’s target in the block model for program comprehension to ease comparison with previous studies.

In Q1 (atom-text), students have 86% success rate. Indented line in Table 1 presents common missed variable names which is different from how those variables were actually used. Four students missed a variable name that was assigned in a *for*-statement and seven students missed a variable name that they used for a reference to a list structure. We did not identify other potential patterns.

The Q2 (relations-execution) was asked from 48 students less than the other two question types as the question could not be generated for programs that did not include a try-except structure. The students who received the question have an 85% success rate. The majority (33) of incorrect answers selected the line that declares the targeted try-block and in 24 cases of those, the correct line is the next line after the selected try-line.

Table 1: The three QLCs, descriptions of their correct and incorrect answer options, and numbers of students who answered each option. Success rates are in bold text and indented lines list student counts for actual labels used as answer options.

Q1 Which of the following are variable names in the program? ($N = 291$)		Students	
• Correct options selected		249	86%
• Missed a variable		22	8%
i (6) list (4) sum (2) <i>one-off names</i> (11)			
• Selected a built-in function		21	7%
print (7) float (7) input (6) len (1)			
• Selected a reserved word		20	7%
return (6) if (5) while (3) try (2)			
except (2) pass (1) for (1)			
• Selected an unused word		3	1%
n (1) total (1) other (1)			
Q2 From which line can program execution jump to line X? ($N = 243$)		Students	
<i>The line in question starts an except-block.</i>		243	
• Correct line		207	85%
int/float (196) input (7) division by zero (4)			
• The related try line		33	14%
• Outside and before the related try-block		3	1%
Q3 Which of the following best describes the purpose of line X? ($N = 291$)		Students	
<i>The line in question calls built-in function input.</i>		234	
• Correct: Accepts new data		227	97%
• Ignores negative input		4	2%
• Is a condition for ending the program		2	1%
• Guards against division by zero		1	0%
<i>The line in question is a condition that skips a block where zero would divide.</i>		57	
• Correct: Guards against division by zero		54	95%
• Is a condition for ending the program		2	4%
• Ignores negative input		1	2%

Table 2: Median course points for students who answered different QLCs correctly (T) or incorrectly (F), numbers of students, degrees of freedom, Mann-Whitney *U*-test statistic, *p*-value for the null hypothesis that the two medians are equal, and common language effect size.

QLC	MD_T	MD_F	n_T	n_F	df	U	p	CLES
Q1	5272	5603.5	249	42	289	4932	.556	.47
Q2	5366	4227.5	207	36	241	4716	.011	.63
Q3	5348	3941	281	10	289	2124	.006	.76

In Q3 (atom-purpose), students have 96% success rate. We did confirm that none of the incorrect answers could be defended with

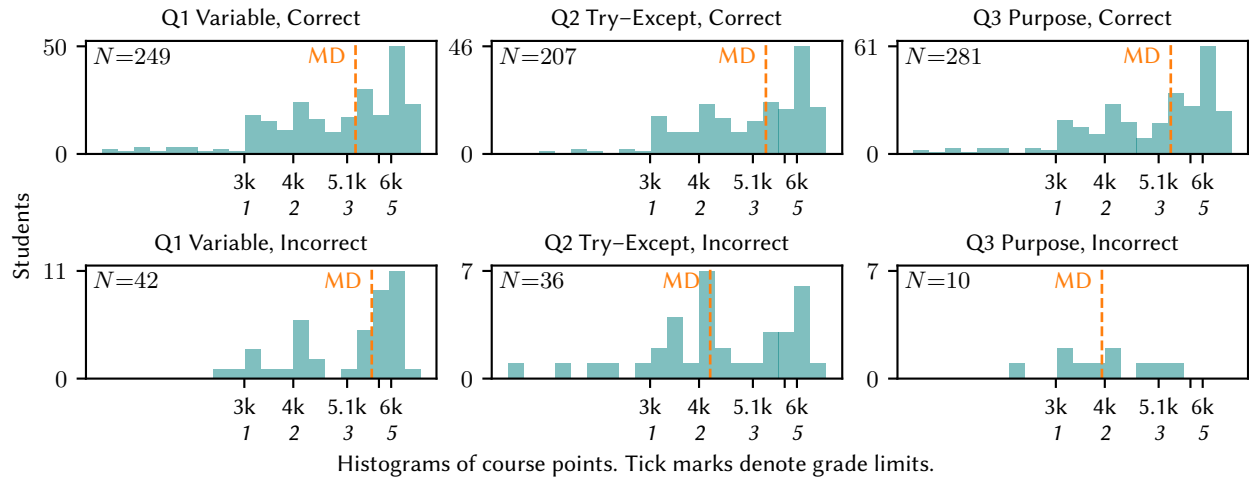


Figure 2: Students' course performance in groups by whether they answered correctly or incorrectly to the QLC types.

true arguments. Possible reasons include selecting “Ignores negative input” when the line has a condition to skip a block where division by zero would happen, and selecting “Is a condition for ending the program” when actually the line accepted new input using a prompt that mentioned how to end the program.

Among the students who submitted the program and answered the QLCs there were 13 cases where the student's final program failed at least one of the functional tests. From these students 100% answered correctly to Q1 and 92% to Q3. 8 of these students received Q2 and 63% answered correctly. The overlap with students who did not fix their program to pass all tests and students who answered QLCs incorrectly is small although students with incomplete programs have relatively less generations and more incorrect answers to Q2.

In Figure 2, we study course points distributions separately for the three different types of QLCs generated and whether the student answered correctly or incorrectly. Our statistical tests reported in Table 2 reveal that students who answered Q2 or Q3 incorrectly had statistically significantly lower median course points than students who answered correctly ($p < .017$) while answers to Q1 divided students into groups having statistically identical means. We calculated medium effect size for Q2 (.63) and large effect size for Q3 (.76). The common language effect size is the probability that in a randomly selected pair the student who answered incorrectly has lower course points than the student who answered correctly.

5 DISCUSSION

5.1 Interpretation and Implications

In this study, we have introduced QLCs incorporated into an automatic assessment system that takes students' program code as input, generates multiple-choice questions that target the program, collects students' answers, and presents automated feedback. We have analyzed the results. In the following, we discuss the findings in the light of the two research questions and compare them with previous studies.

5.1.1 RQ1: How well do students answer QLCs for an exercise targeting prerequisite knowledge? For Q1 and Q2 about 15% of the students had answered incorrectly. This is aligned with previous QLCs that have targeted block model's *text* level or simple cases of *execution*. For these types of QLCs students have typically had above 70% success rates on their first programming course and in many cases 10–15% of answers were incorrect. Lower success rates have been recorded when students have to trace, that is, mentally execute the program to answer QLCs.

Only 5% of students failed Q3 which is comparable to the best success rates recorded in previous studies. There is only one earlier evaluation of *purpose* level QLC where students had 91% success rate [8].

It is surprising that the proportional number of students answering simple questions incorrectly remains the same also after completing their first programming course. Even though the student population is different in these studies, there seems to be a trend that some students can achieve very good program comprehension already during the first couple of weeks in a first programming course, while some students still struggle after completing a whole course. The reasons must be manifold. Yet, we can use the results to explore students' programming knowledge and program comprehension.

In Q1, some students ignored variables that were used for lists or iteration and the difference from assigning a single value to a variable may have affected their thoughts. In Q2, several students may have intentionally selected the line just before where the error occurs which may be related to their mental model of execution. A common reason for failing one or two unit tests for the program are problems in handling unexpected input with a *try*-statement. Those students received less related Q2 and struggled more to answer. Many incorrect answers to Q3 suggest that the student searched for superficial hints in the code line, such as strings of text or numerical values, without understanding the actual purpose of the line. We argue that above are signs of fragile knowledge.

Plagiarism could increase the number of incorrect answers to QLCs as students have weaker understanding of the code they submit in comparison to students who created their own programs. Similarly, they may copy parts of code from examples, online sources and discussions, or occasionally directly from instructors. Given that our programming task is relatively simple and programs are small students could answer QLCs correctly also when they did not create the targeted program themselves. In fact, we expect that students who mastered CS1 should be able to do so.

Currently, artificial intelligence (AI) is readily available for students to produce new code they may not understand but which solves their programming assignment⁵. In some cases AI has been found to write better programs for typical tasks in CS1 than most students [2]. Raman and Kumar [14] argue that, due to advances in AI, changes are required to the focus, pedagogy, and assessment of program writing. They recommend greater emphasis on program testing and comprehension. We believe QLCs have potential to be part of the change.

5.1.2 RQ2: How do the QLC results relate to success on second programming course? All of the researched students did submit a working rainfall-program although 13 cases still failed some functional tests. The students who then answered incorrectly to Q2 or Q3 on average had lower course success. The Q2 and Q3 were designed to target student's program comprehension at block model's execution and purpose levels respectively. Student's fragile knowledge in these areas would be a hindering factor in learning more programming content. We offer this as potential hypothesis for the lower course success. As Q1 is a checkbox-style multiple choice question it required the students' carefully consider every answer option separately. The incorrect options listed in Table 1 suggest that there are probably multiple different reasons for the incorrect answers.

Students in the previous study [16] self-reported terminology as the topic they learned the most from answering QLCs. We assume students can reason about program execution without knowing the terminology that well. We do recognize the value of terminology for communication and continued learning. However, in our study, incorrect interpretations of execution or purpose are more related with lower course success than potential gaps in knowledge of terminology.

5.2 Threats to Validity

As students volunteered for the study there is a risk of selection bias. Of all the students who started the course, 94% gave consent to use their exercise data without any incentives. Unfortunately, 33 students did not complete the studied task. This likely excluded lower achieving students - a possible selection bias.

It is important to note that there are many possible reasons to answer incorrectly to the QLCs that we generated. We assume that these simple questions can be answered correctly without significant extra effort if the code is already understood while producing it. However, there are reasons other than poor understanding, such as carelessness which can make the student select an incorrect option. Poor understanding in turn can be a result of fragile prerequisite

knowledge or even plagiarised code with no intention to understand the code. Some students may have also sought unsupervised help to answer their personal QLCs online.

We chose to have only three simple questions to keep the study short to ensure the participants stayed motivated. The students' programs were manually examined to ensure there were no valid arguments to choose an option designed to be incorrect. However, we did not interview the students to discover why they answered as they did.

In our analysis we use non-parametric statistical tests which are less powerful than *t*-test for normal distributions, thus failure to reject hypothesis of equal means is more likely. In contrast, we test three different variables which increases the chance of finding a case where the test rejects hypothesis of equal means. As a countermeasure for rejecting the true hypothesis we use Bonferroni correction. In the end, the tests did identify two cases of different means from one case of identical means. We note that the population who answered incorrectly to Q3 is small and recommend validation using a larger population.

5.3 Conclusion

In this study, we have confirmed the results of earlier studies [7, 8, 16]. We found similar levels of incorrect answers to simple QLCs on the second programming course as had previously been found for the first weeks of starting programming. Success in two of our QLCs predicted higher median of course points which supports previous findings and adds to the testimony that QLCs can detect factors that are relevant for learning. Thus, we argue QLCs are a promising new question type to give early warnings on weak prerequisite knowledge, which can be taken into account in the course design. Some students might benefit from a more comprehensive recap round in the beginning of the second programming course. QLCs could complement concept inventories to provide a wider set of problems to reveal students' misconceptions that hinder learning.

More generally, we conclude that QLCs are a promising extension to automated assessment systems. QLCs can be utilized to extend the assessment of student-created programs beyond unit testing and style linting to program comprehension [9]. If implemented well, the shift in assessment may also affect student's attitudes on the importance of understanding programs [15].

The reasons why students struggle to interpret their own code are not known. Novice coders tend to make extensive use of programming forums [23] such as Stack Overflow and might use code snippets that contain code they might not fully understand. The recent introduction of free and easily available AI-based code generation tools is likely to exacerbate the problem further. As students start using such tools while learning programming, the teaching should focus more on tools and effective strategies for program comprehension.

In the future, more studies are needed to analyse different types of QLCs and why some students struggle to answer them correctly. Moreover, with AI-assisted programming on the rise and student enrollments high, it is ever more timely to find automated triggers for self-explanations of code and self-regulation. QLCs hold the potential to achieve that, but further research is needed to tap into that potential.

⁵For example Github Copilot and ChatGPT are AI tools that employ successful large language models to generate program code.

REFERENCES

- [1] Kathryn Cunningham, Rahul Agrawal Bejarano, Mark Guzdial, and Barbara Ericson. 2020. "I'm Not a Computer": How Identity Informs Value and Expectancy During a Programming Activity. In *The Interdisciplinarity of the Learning Sciences, 14th International Conference of the Learning Sciences (Nashville, Tennessee) (ICLS 2020, Vol. 2)*. International Society of the Learning Sciences (ISLS), USA, 705–708. <https://repository.isls.org/handle/1/6733>
- [2] James Finnie-Ansley, Paul Denny, Brett A. Becker, Andrew Luxton-Reilly, and James Prather. 2022. The Robots Are Coming: Exploring the Implications of OpenAI Codex on Introductory Programming. In *Proceedings of the 24th Australasian Computing Education Conference (Virtual Event, Australia) (ACE '22)*. Association for Computing Machinery, New York, NY, USA, 10–19. <https://doi.org/10.1145/3511861.3511863>
- [3] Cruz Izu and Claudio Mirolo. 2020. Comparing Small Programs for Equivalence: A Code Comprehension Task for Novice Programmers. In *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education (Trondheim, Norway) (ITiCSE '20)*. Association for Computing Machinery, New York, NY, USA, 466–472. <https://doi.org/10.1145/3341525.3387425>
- [4] Cruz Izu, Carsten Schulte, Ashish Aggarwal, Quintin Cutts, Rodrigo Duran, Mirela Gutica, Birte Heinemann, Eileen Kraemer, Violetta Lonati, Claudio Mirolo, and Renske Weeda. 2019. Fostering Program Comprehension in Novice Programmers - Learning Activities and Learning Trajectories. In *Proceedings of the Working Group Reports on Innovation and Technology in Computer Science Education (Aberdeen, Scotland UK) (ITiCSE-WGR '19)*. Association for Computing Machinery, New York, NY, USA, 27–52. <https://doi.org/10.1145/3344429.3372501>
- [5] Cazembe Kennedy and Eileen T. Kraemer. 2019. Qualitative Observations of Student Reasoning: Coding in the Wild. In *Proceedings of the 2019 ACM Conference on Innovation and Technology in Computer Science Education (Aberdeen, Scotland UK) (ITiCSE '19)*. Association for Computing Machinery, New York, NY, USA, 224–230. <https://doi.org/10.1145/3304221.3319751>
- [6] Viraj Kumar. 2021. Refute: An Alternative to 'Explain in Plain English' Questions. In *Proceedings of the 17th ACM Conference on International Computing Education Research (Virtual Event, USA) (ICER 2021)*. Association for Computing Machinery, New York, NY, USA, 438–440. <https://doi.org/10.1145/3446871.3469791>
- [7] Teemu Lehtinen, Lassi Haaranen, and Juho Leinonen. 2023. Automated Questionnaires About Students' JavaScript Programs: Towards Gauging Novice Programming Processes. In *Proceedings of the 25th Australasian Computing Education Conference (Melbourne, VIC, Australia) (ACE '23)*. Association for Computing Machinery, New York, NY, USA, 49–58. <https://doi.org/10.1145/3576123.3576129>
- [8] Teemu Lehtinen, Aleksu Lukkarinen, and Lassi Haaranen. 2021. Students Struggle to Explain Their Own Program Code. In *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education V. 1 (Virtual Event, Germany) (ITiCSE '21)*. Association for Computing Machinery, New York, NY, USA, 206–212. <https://doi.org/10.1145/3430665.3456322>
- [9] Teemu Lehtinen, André L. Santos, and Juha Sorva. 2021. Let's Ask Students About Their Programs, Automatically. In *2021 IEEE/ACM 29th International Conference on Program Comprehension (ICPC)*. IEEE, New York, NY, USA, 467–475. <https://doi.org/10.1109/ICPC52881.2021.00054>
- [10] Raymond Lister, Elizabeth S. Adams, Sue Fitzgerald, William Fone, John Hamer, Morten Lindholm, Robert McCartney, Jan Erik Moström, Kate Sanders, Otto Seppälä, Beth Simon, and Lynda Thomas. 2004. A Multi-National Study of Reading and Tracing Skills in Novice Programmers. In *Working Group Reports from ITiCSE on Innovation and Technology in Computer Science Education (Leeds, United Kingdom) (ITiCSE-WGR '04)*. Association for Computing Machinery, New York, NY, USA, 119–150. <https://doi.org/10.1145/1044550.1041673>
- [11] Sandra Madison and James Gifford. 2002. Modular Programming. *Journal of Research on Technology in Education* 34, 3 (2002), 217–229. <https://doi.org/10.1080/15391523.2002.10782346>
- [12] Claudio Mirolo, Cruz Izu, and Emanuele Scapin. 2020. High-School Students' Mastery of Basic Flow-Control Constructs through the Lens of Reversibility. In *Proceedings of the 15th Workshop on Primary and Secondary Computing Education (Virtual Event, Germany) (WiPSCE '20)*. Association for Computing Machinery, New York, NY, USA, Article 15, 10 pages. <https://doi.org/10.1145/3421590.3421603>
- [13] D. N. Perkins and Fay Martin. 1986. Fragile Knowledge and Neglected Strategies in Novice Programmers. In *Papers Presented at the First Workshop on Empirical Studies of Programmers* (Washington, D.C., USA). Ablex Publishing Corp., USA, 213–229.
- [14] Arun Raman and Viraj Kumar. 2022. Programming Pedagogy and Assessment in the Era of AI/ML: A Position Paper. In *Proceedings of the 15th Annual ACM India Compute Conference (Jaipur, India) (COMPUTE '22)*. Association for Computing Machinery, New York, NY, USA, 29–34. <https://doi.org/10.1145/3561833.3561843>
- [15] Kay Sambell and Liz McDowell. 1998. The Construction of the Hidden Curriculum: messages and meanings in the assessment of student learning. *Assessment & Evaluation in Higher Education* 23, 4 (1998), 391–402. <https://doi.org/10.1080/0260293980230406>
- [16] André Santos, Tiago Soares, Nuno Garrido, and Teemu Lehtinen. 2022. Jask: Generation of Questions About Learners' Code in Java. In *Proceedings of the 27th ACM Conference on Innovation and Technology in Computer Science Education Vol. 1 (Dublin, Ireland) (ITiCSE '22)*. Association for Computing Machinery, New York, NY, USA, 117–123. <https://doi.org/10.1145/3502718.3524761>
- [17] Carsten Schulte. 2008. Block Model: An Educational Model of Program Comprehension as a Tool for a Scholarly Approach to Teaching. In *Proceedings of the Fourth International Workshop on Computing Education Research (Sydney, Australia) (ICER '08)*. Association for Computing Machinery, New York, NY, USA, 149–160. <https://doi.org/10.1145/1404520.1404535>
- [18] Carsten Schulte, Tony Clear, Ahmad Taherkhani, Teresa Busjahn, and James H. Paterson. 2010. An Introduction to Program Comprehension for Computer Science Educators. In *Proceedings of the 2010 ITiCSE Working Group Reports (Ankara, Turkey) (ITiCSE-WGR '10)*. Association for Computing Machinery, New York, NY, USA, 65–86. <https://doi.org/10.1145/1971681.1971687>
- [19] Amal A. Shargabi, Syed Ahmad Aljunid, Muthukkaruppan Annamalai, and Abdullah Mohd Zin. 2020. Performing Tasks Can Improve Program Comprehension Mental Model of Novice Developers: An Empirical Approach. In *Proceedings of the 28th International Conference on Program Comprehension (Seoul, Republic of Korea) (ICPC '20)*. Association for Computing Machinery, New York, NY, USA, 263–273. <https://doi.org/10.1145/3387904.3389277>
- [20] Simon. 2011. Assignment and Sequence: Why Some Students Can't Recognise a Simple Swap. In *Proceedings of the 11th Koli Calling International Conference on Computing Education Research (Koli, Finland) (Koli Calling '11)*. Association for Computing Machinery, New York, NY, USA, 10–15. <https://doi.org/10.1145/2094131.2094134>
- [21] Elliot Soloway, Kate Ehrlich, Jeffrey Bonar, and Judith Greenspan. 1982. What do novices know about programming? *Directions in Human-Computer Interaction* 6 (1982), 27–54.
- [22] Dieter Vogts. 2009. Plagiarising of Source Code by Novice Programmers a "Cry for Help?". In *Proceedings of the 2009 Annual Research Conference of the South African Institute of Computer Scientists and Information Technologists (Vanderbijlpark, Emfuleni, South Africa) (SAICSIT '09)*. Association for Computing Machinery, New York, NY, USA, 141–149. <https://doi.org/10.1145/1632149.1632168>
- [23] David Wong-Aitken, Diana Cukierman, and Parmit K. Chilana. 2022. "It Depends on Whether or Not I'm Lucky" How Students in an Introductory Programming Course Discover, Select, and Assess the Utility of Web-Based Resources. In *Proceedings of the 27th ACM Conference on Innovation and Technology in Computer Science Education Vol. 1 (Dublin, Ireland) (ITiCSE '22)*. Association for Computing Machinery, New York, NY, USA, 512–518. <https://doi.org/10.1145/3502718.3524751>