Analyzing the uses and perceptions of computer science students towards generative AI tools

Alisha HASAN[1], Bogdan SIMION[1], Florin POP[2,3,4]

[1] University of Toronto Mississauga,
Mississauga Road, Mississauga, Ontario, Canada

[2] Politehnica Bucharest National University of Science and Technology,
Splaiul Independentei 313, Bucharest, Romania

[3] National Institute for Research & Development in Informatics – ICI Bucharest,
8-10, Maresal Averescu Avenue, 1st District, Bucharest, Romania

[4] Academy of Romanian Scientists, 3 Ilfov Street, Bucharest, Romania

alisha.hasan@mail.utoronto.ca, bogdan@cs.toronto.edu,
florin.pop@upb.ro, florin.pop@ici.ro

Abstract: *With the evolution of generative AI (GenAI) tools based on Large Language Models (LLMs), stakeholders in Computer Science education have sought to leverage opportunities and mitigate risks that these tools may entail in the educational process. Thus, it is crucial to understand students' uses and perceptions of GenAI tools in their learning. Our work captures how computer science students use GenAI tools that produce code or text. Additionally, we investigate student perceptions in terms of important factors such as perceived usefulness, correctness, reliability, accessibility, trust in responses, motivation to learn, confidence building, potential dangers, and perceived implications on the software development industry. After surveying students from three institutions, we conducted a mixed-methods analysis and noted that student uses of GenAI tools are typically educational, with few students using them for purposes detrimental to their learning such as generating solutions. Additionally, students are more aware of the potential for misleading incorrect responses from GenAI tools and the pitfalls of over-reliance on such tools than educators might expect. Overall, our findings are useful for computer science educators to be mindful of students' uses and perspectives on GenAI, to better guide students towards positive learning outcomes and foster good learning habits.*

Keywords: Generative AI, LLMs, CS Education, Student Perceptions

1. Introduction

The recent evolution of Large Language Models (LLMs) has had a significant impact on several areas of activity, but especially in education. These tools are gradually able to solve programming exercises to a greater and greater extent, especially when carefully crafting prompts (Denny et al., 2023; Finnie-Ainsley et al., 2023; Wermelinger, 2023; Vahid et al., 2024). As instructor attitudes towards LLMs evolved (Lau & Guo, 2023; Sheard et al., 2024) instructors

https://doi.org/10.58503/icvl-v19y202430

acknowledged that there are both learning opportunities and threats that these tools can pose in the educational process. Opportunities include automatic generation of contextualized exercises, (Sarsa et al., 2022; del Carpio Gutierrez et al., 2024) autograding and automated virtual Teaching Assistants (Liu et al., 2024; Liu & M'hiri, 2024), aiding students in debugging and understanding compiler error messages, (Al-Hossaimi et al., 2024; Taylor et al., 2024; Wang et al., 2024) as well as enhancing students' prompt engineering skills via "prompt problems" (Denny et al., 2024).

While the literature has focused largely on applications of LLMs or instructor attitudes, very few studies (Amoozadeh et al., 2024) focused on how students use these tools or their perceptions towards these tools as learning supports or their pitfalls. This work aims to analyze student uses and perceptions of such tools in the current computational landscape where these tools have evolved at a rapid pace. We investigate the following research questions:

- To what extent do students use such tools to generate text or code, for either schoolwork, or personal projects?
- What are the perceptions of students with respect to GenAI tools? Specifically, we look at factors such as perceived usefulness, correctness, reliability, accessibility, trust in the results, motivation to learn, confidence building, dangers that these tools are perceived to pose, and potential implications on software development jobs.

We surveyed students from three teaching and research institutions for higher learning and analyzed the 147 responses, comprising both undergraduate and graduate students. In the following sections, we discuss prior work related to this topic, then our data collection and methodology, followed by our results and discussing our findings and potential limitations, before concluding in section 7.

2. Related work

2.1 Use Cases of GenAI tools by computing students

The developments and increased usage of GenAI tools, spurred research on how to leverage these tools within the programming community, both from the teaching and learning perspectives. While most studies have focused on instructor uses and applications of LLMs in the educational process, a few studies have explored the different ways in which GenAI tools are used by programming students. One study found that students typically turn to GenAI in a manner conducive to learning, e.g., as a specialized search engine, to better grasp new concepts, to study for exams, or to get guidance, rather than generate solutions to assigned work (Rogers et al., 2024). Another study found that students typically turn to GenAI to seek help and understand code rather than generate new code (Amoozadeh et al., 2024).

2.2 Computing students' perceptions of GenAI

While studies researching student perceptions of GenAI tools in Computer Science are scarce, some empirical evidence shows students expressing positive sentiments towards GenAI use for learning purposes and idea generation both in computer science (Amoozadeh et al., 2024) and in the context of other disciplines (Arowosegbe et al., 2024; Chan & Hu, 2023). In a study focused on analyzing trust in GenAI tools (Amoozadeh et al., 2024), computing students expressed a neutral perception towards GenAI tools. While most expressed a certain level of lack of trust in the generated output, an equally significant portion of students regarded GenAI as generally helpful due to the ability to condense large amounts of information and provide help. Other studies outside of Computer Science illustrated student concerns regarding increased usage of LLMs correlating to a higher level of overreliance (Chan & Hu, 2023; Salifu et al., 2024). Additionally, research studies highlighted the negative perceptions towards GenAI due to perceived career threats, misuse for academic dishonesty, and privacy issues (Arowosegbe et al., 2024; Kamoun et al. 2024).

3. Data collection and methodology

3.1 Survey and demographics

We conducted a survey at three Eastern European research and teaching institutions for higher learning, with a mix of undergraduate and graduate students. The survey was distributed on all the undergraduate and graduate student mailing lists and participation was not incentivized with any form of compensation.

Out of 147 responders, we had 85 undergraduate students fairly evenly distributed across the four years of study, 40 Master's students, 19 PhD students, and 3 postdoctoral fellows. The demographics included 116 male and 24 female responders, with the rest opting to not disclose this information; most students (93.9%) were of European descent, with the rest of West Central Asian and Middle Eastern, South Asian, African, or Carribbean backgrounds, and 6 responders preferring not to answer. Although most applicants did not have English as their native language, their command of English was solid as this is a curricular requirement at the respective institutions.

The responders' focuses or specializations include a mix of areas of computer science with most students specializing in a combination of Computer Systems (including Distributed Systems, Operating Systems, Compilers and Databases - 51.7%), Artificial Intelligence and Machine Learning (27.9%), Web Development (27.2%), Computational Theory (26.5%), Computer Architecture and Embedded Systems (22%), Computer Security (18.3%), Computer Graphics (8.8%), and HCI (2.7%). The percentages do not add up to 100% as each student typically specializes in multiple areas. On average, 22.7% of students were the first generation to go to university in their family, while 71.4% had at least one parent who had already completed a graduate or postgraduate degree, and the rest chose not to respond.

Aside from questions to gather demographic data, we asked a set of Likert-scale questions as well as open questions, discussed further in the next sections.

3.2 Likert-scale survey questions

The Likert-scale questions were grouped into two sets of questions, to capture:

1. Frequency and types of usage of GenAI, and
2. Perceptions of students towards GenAI tools, including perceived usefulness, correctness, reliability, accessibility, trust in the results, motivation to learn, confidence building, dangers that these tools are perceived to pose, and potential implications on software development jobs.

The set of Likert-scale questions on student usage types, as well as those related to student perceptions are presented in the Results section in Figures 1 and 2, to visualize the questions closer to the response distributions. By asking a wide variety of questions in different forms, we aimed to capture a variety of responses from undergraduate and graduate students of varying academic levels.

3.3 Open-ended questions and coding methodology

Aside from the Likert-scale questions, we asked several open-ended questions (included in Table 1). Additionally, we had one final question to collect anything else that students wanted to express regarding their views on GenAI tools or uses. To analyze the responses to these open questions, we adopted an open thematic coding process to extract common themes in student perceptions. Each response was coded into one or more themes, then all responses were reviewed to ensure consistency.

Table 1. Open-ended Questions, Emerging Themes, and Number of Responses

QUESTION	ASSOCIATED THEMES (MENTIONS)	# OF RESPONSES
1. I have used generative AI tools to generate text or explanations for other purposes than in the questions above (please specify).	Learning Purposes **(10)** Does Not Use **(7)** Writing **(6)** Entertainment **(5)** Idea Generation **(5)** Programming **(5)** Planning/Organization **(5)** Information Retrieval **(3)** Time Saving **(2)**	51
2. I have often used generative AI tools to generate code for other purposes than in the questions above (please specify).	Does Not Use **(16)** Code Generation **(7)** Testing/Debugging **(4)** Optimization/Time Saving **(2)**	35

QUESTION	ASSOCIATED THEMES (MENTIONS)	# OF RESPONSES
3. Elaborate on what you think are the benefits or opportunities of generative AI tools that can generate text or code.	Time Saving **(43)** Boosted Productivity **(15)** Learning **(15)** Automation **(10)**	85
4. Elaborate on what you think are the risks or dangers of generative AI tools that can generate text or code.	Overreliance **(37)** Unreliable **(27)** Privacy Concerns **(7)** Job Displacement **(4)** Unsafe **(3)** Biased **(2)**	85

4. Results

4.1 Student uses of GenAI tools

Our data, as shown in Table 1, shows that most students have used GenAI tools, with 64% of students using them often or very frequently, and ~24% of students stating they sometimes use such tools. Additionally, around 58% of responders use GenAI tools often or very frequently when learning new concepts, while only a little over 5% never use GenAI for such purposes. In terms of writing tasks (such as reports, essays or documents), the student responses fall mostly in the middle when it comes to generating full write-ups or reformulating a piece of writing.

For coding tasks, most students have responded Sometimes, Often, or Very Frequently when asked whether they used GenAI tools to generate code. The percentages for this question align with their responses on how much they use GenAI for helper code, which implies that the main purpose to using GenAI is as an auxiliary tool. Additionally, around 55% of students never used it for detrimental purposes such as generating entire solutions, and only 2.8% frequently resort to such practices. This is mirrored by the responses to questions on generating code for personal projects.

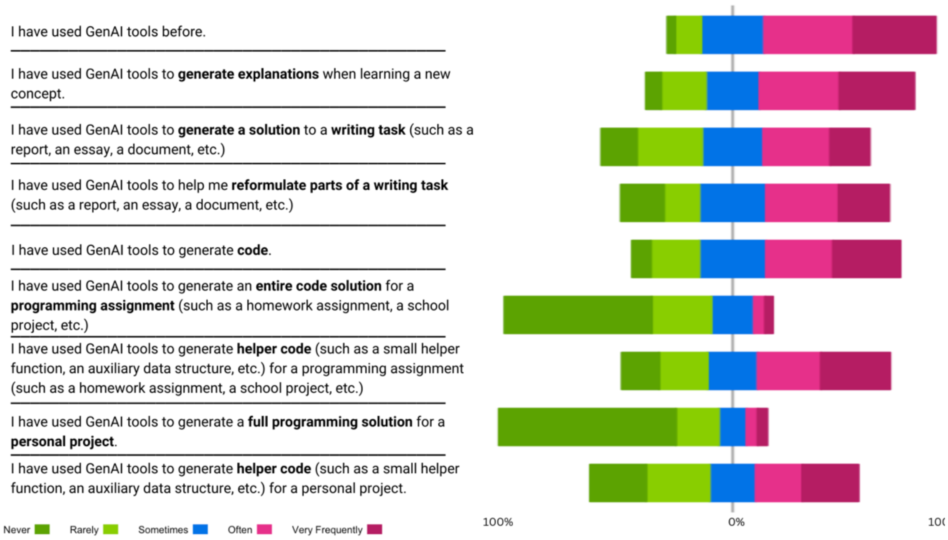

Figure 1. Student Responses on Their Uses of GenAI Tools.

All in all, this suggests that overall, most students use GenAI tools for legitimate learning purposes (e.g., learning new concepts, revising written work, or helper code), while very few resorts to undesirable learning practices. This is further supported by the themes we captured from the open questions (summarized in Table 1), where 52% of students reported GenAI tools as beneficial for time saving purposes, for boosting overall productivity when completing projects, seamless information retrieval, tedious task completion, and streamlining overall learning by acting as a personalized tutor. It is worth mentioning that there were no statistically significant differences across demographics, other than a slightly higher percent of graduate students using GenAI to generate code, compared to undergraduate students, albeit not statistically significant.

4.2 Student perceptions of generative AI

In terms of the ability of GenAI tools to produce *correct* explanations, text, or code, the majority leaned towards either Agree (45-47%) or felt neutral (30-35%), hinting towards cautious optimism mixed with some skepticism about the output correctness, and which was even more prominent in terms of *completeness* of explanations, text, or code. This outlook can be explained in part by the responses on whether *prompts and follow-ups* are necessary to get correct results, to which the vast majority (83-84%) agreed were required.

Most responders (85-92%) agreed on *accessibility* of explanations, text, and code generated by GenAI tools, which is encouraging to see that the output is generally not beyond a student's ability to grasp. Most students agreed on the ability of such tools to *improve conceptual knowledge*, but a slightly higher number of students had a more neutral view on GenAI's ability to *improve programming skills*. The responses were also mixed in terms of GenAI tools being able to *increase motivation to learn*, *confidence in one's abilities*, their *reliability*, or *trustworthiness*. However, two thirds of the responders had an *overall positive view* of the current state of GenAI tools, and 87% of responders believed that these tools will *improve over time*.

The *risks* of GenAI were also clear to most responders, with 73% of responders agreeing on *potential dangers* of GenAI. Nevertheless, most agreed that these tools will become part of the *software development cycle*, yet not fully replacing software development jobs.

The thematic analysis of the open questions enhances these findings. Many responders found GenAI tools to be overall useful when asked about any additional thoughts in the last survey question. However, it is notable that most students specified that GenAI is useful when used correctly, and a major ethical concern when used improperly. It is noteworthy that students reported concerns regarding the rise of GenAI tools despite continuing to use them. For example, the themes from Table 1 highlight student beliefs that continued use of GenAI will result in overreliance and obsolescence of skills that GenAI tools replace. This is further supported in the open questions where many expressed concerns surrounding loss of skill due to GenAI use. The most frequently cited concern was the potential deterioration of writing skills as a direct result of relying on GenAI to perform writing tasks. It is additionally notable that a few students expressed either fear towards a future with GenAI, or excitement about the possibilities of GenAI boosting productivity.

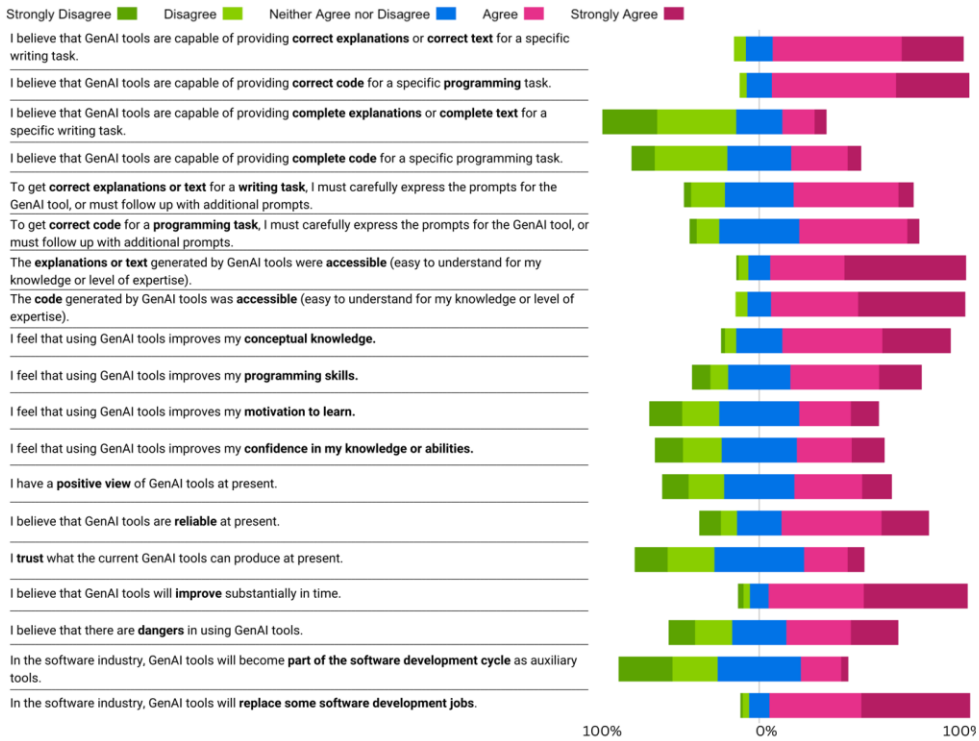

Figure 2. Student Responses to Questions on Their Perceptions Towards GenAI Tools

5. Discussion

Student responses illustrate an overall positive view of GenAI tools and a perception of GenAI tools being helpful in the learning process. For example, some students reported using GenAI to generate explanations of concepts at varying levels of intricacy and expressed satisfaction with the results, which has encouraged their continued usage for learning purposes. It appears that our responders tend to integrate GenAI as a tool to aid their learning, unlike other studies which have found that students viewed GenAI tools as a means of academic dishonesty (Arowosegbe et al., 2024; Black, 2024; Kamoun et. al., 2024). It is important to recognize though that the role GenAI plays in student education can heavily revolve around the way it has been presented to students by faculty. If instructors aim to integrate GenAI in a productive manner, positive sentiments towards the effect it has on student learning may be more likely.

Despite potential benefits of integrating GenAI tools in student curriculums, it is important to acknowledge the concerns that students expressed about long-term effects of using GenAI tools. For instance, habitual use due to convenience may lead to overreliance on such tools. According to our survey results, while it is widely perceived that GenAI tools can enhance skill development, cautious GenAI use is required to keep it a facilitating tool rather than a skill replacement.

Educators can reinforce to students that this balance is essential to mitigate negative impacts on student skill development.

A notable takeaway of the perceived benefits and risks among students is that, paradoxically, despite the students' primary concern being overreliance on GenAI tools, this concern did not significantly deter students from using these tools, nor impact their overall positive view on such tools. Potential reasons include benefits outweighing any concerns of degraded skills or may reflect a mature attitude of our responder sample being able to recognize where to draw the line in using such tools.

6. Limitations

Although our data comes from several institutions and students with diverse levels of study and specializations, it is limited to one geographical context. While our findings align with the related literature, we acknowledge this limitation to generality, and call for replication studies at institutions in other geographical locations, with different demographics and curricular approaches.

7. Conclusion

In this paper, we explored the perceptions of GenAI among students at three Eastern European research and teaching institutions. Our results indicate that students use these tools for legitimate learning purposes and are aware of the risks associated with GenAI. Overall, students were most vocal about the concern of GenAI fostering a general sense of overreliance among users that leads to the skill deterioration; on the other hand, the most recognized benefit is using GenAI as an explanatory tool.

Understanding student views on the potential benefits and risks of GenAI helps us learn how to carefully support students into developing good learning habits and avoiding GenAI pitfalls. As future directions, it is important to use such findings to develop best practices for integrating GenAI in course policies and pedagogic guidelines, to maximize benefits while mitigating risks. Additionally, although beyond the scope of our work, an open question remains on the actual long-term impact of GenAI on student skills and overall learning.

REFERENCES

Al-Hossaimi, E., Bunescu. R., Smith, J. & Teehan, J. (2024) Can Language Models Employ the Socratic Method? Experiments with Code Debugging (2024). *SIGCSE Technical Symposium 2024*. pp. 53-59. doi:10.1145/3626252.3630799.

Amoozadeh, A., Daniels, D., Nam, D., Kumar, A., Chen, S., Hilton, M., Ragavan, S. & Alipour, M. (2024) Trust in Generative AI among Students. *SIGCSE Technical Symposium 2024*. pp. 67-73. doi:10.1145/3626252.3630842.

Arowosegbe, A., Alqahtani, J. S. & Oyelade, T. (2024) Students' Perception of Generative AI Use for Academic Purpose in UK Higher Education. *Preprints.org*, https://www.preprints.org/manuscript/202405.1158/v1 [Accessed 20th August 2024].

Black, E. S. (2024) Working Smarter: A Quantitative Investigation Into Higher Education Faculty's Perceptions, Adoption, and Use of Generative Artificial Intelligence (AI) in Alignment With the Learning Sciences and Universal Design for Learning. *Drexel University ProQuest Dissertations & Theses*, https://www.proquest.com/docview/3077461163/fulltextPDF?sourcetype=Dissertations%20&%20Theses.

Chan, C. K. Y. & Hu, W. (2023) Students' voices on generative AI: perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education.* 20(1). doi:10.1186/s41239-023-00411-8.

del Carpio Gutierrez, A., Denny, P. & Luxton-Reilly, A. (2024) Evaluating Automatically Generated Contextualised Programming Exercises. *SIGCSE Technical Symposium 2024*. pp. 289-295. doi:10.1145/3626252.3630863.

Denny, P., Kumar, V. & Giacaman, N. (2023) Conversing with Copilot: Exploring Prompt Engineering for Solving CS1 Problems Using Natural Language. *SIGCSE Technical Symposium 2023*. pp. 1136–1142. doi:10.1145/3545945.3569823.

Denny, P., Leinonen, J., Prather, J., Luxton-Reilly, A., Amarouche, T., Becker, B. & Reeves, B. (2024) Prompt Problems: A New Programming Exercise for the Generative AI Era. *SIGCSE Technical Symposium 2024*. pp. 296–302. doi:10.1145/3626252.3630909.

Finnie-Ainsley, J., Denny, P., Becker, B., Luxton-Reilly, A. & Prather, J. (2022) The Robots Are Coming: Exploring the Implications of OpenAI Codex on Introductory Programming. *Australasian Computing Education Conference (ACE) 2022*. pp. 10-19. doi:10.1145/3511861.3511863.

Kamoun, F., El Ayeb, W., Jabri, I., Sifi, S. & Iqbal, F. (2024) Exploring Students' And Faculty's Knowledge, Attitudes, And Perceptions Towards ChatGPT: A Cross-Sectional Empirical Study. *Journal of Information Technology Education: Research*. 23, 1-33. doi:10.28945/5239.

Lau, S. & Guo, P. (2023) From "Ban It Till We Understand It" to "Resistance is Futile": How University Programming Instructors Plan to Adapt as More Students Use AI Code Generation and Explanation Tools such as ChatGPT and GitHub Copilot. *Conference on International Computing Education Research (ICER) 2023*. doi:10.1145/3568813.3600138.

Liu, M. & M'hiri, F. (2024) Beyond Traditional Teaching: Large Language Models as Simulated Teaching Assistants in Computer Science. *SIGCSE Technical Symposium 2024*. pp. 743–749. doi:10.1145/3626252.3630789.

Liu, E. Z., Yuan, D., Ahmed, A., Cornwall, E., Woodrow, J., Burns, K., Nie, A., Brunskill, E., Piech, C. & Finn, C. (2024) A Fast and Accurate Machine Learning

Autograder for the Breakout Assignment. *SIGCSE Technical Symposium 2024.* pp. 736-743. doi:10.1145/3626252.3630759.

Salifu, I., Arthur, F., Arkorful, V., Nortey, S. A. & Osei-Yaw, R. S. (2024) Economics students' behavioural intention and usage of ChatGPT in higher education: a hybrid structural equation modelling-artificial neural network approach. *Cogent Social Sciences.* 10(1). doi:10.1080/23311886.2023.2300177.

Sarsa, S., Denny, P., Hellas, A. & Leinonen, J. (2022) Automatic Generation of Programming Exercises and Code Explanations Using Large Language Models. *Conference on International Computing Education Research (ICER) 2022*. doi:10.1145/3501385.3543957.

Sheard, J., Denny, P., Hellas, A., Leinonen, J. & Malmi, L. (2024) Instructor Perceptions of AI Code Generation Tools: A Multi-Institutional Interview Study. *SIGCSE Technical Symposium 2024.* pp. 1223–1229. doi:10.1145/3626252.3630880.

Taylor, A., Vassar, A., Renzella, J. & Pearce, H. (2024) dcc --help: Transforming the Role of the Compiler by Generating Context-Aware Error Explanations with Large Language Models. *SIGCSE Technical Symposium 2024*. pp: 1314-1320. doi:10.1145/3626252.3630822.

Vahid, F., Areizaga, L. & Pang, A. (2023) ChatGPT and Cheat Detection in CS1 Using a Program Autograding System. *SIGCSE Technical Symposium 2024*. pp. 367-373. doi:10.1145/3649217.3653558.

Wang, S., Mitchell, J. C. & Piech, C. (2024) A Large Scale RCT on Effective Error Messages in CS1. *SIGCSE Technical Symposium 2024.* pp. 1395-1401. doi:10.1145/3626252.3630764.

Wermelinger, M. (2023) Using GitHub Copilot to Solve Simple Programming Problems. *SIGCSE Technical Symposium 2023*. pp. 172-178 doi:10.1145/3545945.3569830.