

Humanizing Automated Programming Feedback: Fine-Tuning Generative Models with Student-Written Feedback

Victor-Alexandru Pădurean
MPI-SWS
vpadurea@mpi-sws.org

Michael Liut
University of Toronto Mississauga
michael.liut@utoronto.ca

Tung Phung
MPI-SWS
mphung@mpi-sws.org

Juho Leinonen
Aalto University
juho.2.leinonen@aalto.fi

Adish Singla
MPI-SWS
adishs@mpi-sws.org

Nachiket Kotalwar
MPI-SWS
nkotalwa@mpi-sws.org

Paul Denny
University of Auckland
paul@cs.auckland.ac.nz

ABSTRACT

The growing need for automated and personalized feedback in programming education has led to recent interest in leveraging generative AI for feedback generation. However, current approaches tend to rely on prompt engineering techniques in which predefined prompts guide the AI to generate feedback. This can result in rigid and constrained responses that fail to accommodate the diverse needs of students and do not reflect the style of human-written feedback from tutors or peers. In this study, we explore learnersourcing as a means to fine-tune language models for generating feedback that is more similar to that written by humans, particularly peer students. Specifically, we asked students to act in the flipped role of a tutor and write feedback on programs containing bugs. We collected approximately 1,900 instances of student-written feedback on multiple programming problems and buggy programs. To establish a baseline for comparison, we analyzed a sample of 300 instances based on correctness, length, and how the bugs are described. Using this data, we fine-tuned open-access generative models, specifically Llama3 and Phi3. Our findings indicate that fine-tuning models on learnersourced data not only produces feedback that better matches the style of feedback written by students, but also improves accuracy compared to feedback generated through prompt engineering alone, even though some student-written feedback is incorrect. This surprising finding highlights the potential of student-centered fine-tuning to improve automated feedback systems in programming education.

Victor-Alexandru Pădurean, Tung Phung, Nachiket Kotalwar, Michael Liut, Juho Leinonen, Paul Denny, and Adish Singla. Humanizing Automated Programming Feedback: Fine-Tuning Generative Models with Student-Written Feedback. In Caitlin Mills, Giora Alexandron, Davide Taibi, Giosuè Lo Bosco, and Luc Paquette (eds.) Proceedings of the 18th International Conference on Educational Data Mining, Palermo, Italy, July, 2025, pp. 434–441. International Educational Data Mining Society (2025).

© 2025 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.
<https://doi.org/10.5281/zenodo.15870290>

Keywords

programming feedback, fine-tuning, generative AI

1. INTRODUCTION

Generative AI offers great potential to enhance programming education by providing personalized feedback to students. This can complement feedback from human tutors by offering continuous, around-the-clock support [6, 41]. However, students often perceive AI-generated feedback as less effective than human-written feedback because it may not fully address their needs or provide the depth of support they expect [28]. In particular, students value the emotional sensitivity of human feedback. Therefore, to maximize student adoption, it is essential to refine AI models to deliver feedback with human-like characteristics, similar to what students would receive from tutors or peers.

Most previous studies have focused on the use of prompt-engineering techniques to customize feedback to meet expert educators' standards [18, 42]. Furthermore, the quality of feedback is typically evaluated via educator-created rubrics [32, 33], emphasizing conciseness and appropriateness (i.e., not explicitly pointing out how to fix bugs). This approach suffers from several limitations. First, it demands extensive effort from experts to engineer prompts which may not generalize across domains. Second, the generated feedback tends to be rigid and constrained (e.g., spanning one or two sentences and not containing code), lacking flexibility for addressing specific needs and different types of bugs. More generally, this approach does not capture student preferences as well as the dynamic and adaptable characteristics of humans, such as tutors and peers.

To address this gap, we leverage learnersourced data for fine-tuning generative models, combining the reliability of student-written feedback with generative models' ability to generate consistent feedback at scale. By fine-tuning on 1,920 student-created feedback instances, we align feedback generated by AI more closely with student communication styles, making it more concise, peer-like, and adaptable, as

I'd be happy to help!

The main issue with your original program is that you're adding the index 'i' to the sum instead of the value at that index 'values[i]'. This is why you're getting the wrong output.

Here's a breakdown of what's happening:

- * When 'i' is 0, 'values[i]' is 10, which is positive, so you add 0 to the sum (because 'i' is 0).
- * When 'i' is 1, 'values[i]' is 20, which is positive, so you add 1 to the sum (because 'i' is 1).
- * When 'i' is 2, 'values[i]' is 30, which is positive, so you add 2 to the sum (because 'i' is 2).

So, the sum becomes $0 + 1 + 2 = 3$, which is why you're getting the output 3.

By changing the line 'sum += i;' to 'sum += values[i];', you're correctly adding the positive values to the sum, which is why the fixed program produces the correct output.

Here are some additional suggestions to improve your code:

- * Consider adding a check to ensure that the input array is not null before trying to access its elements. This can help prevent crashes or unexpected behavior if the input array is invalid.
- * You might want to consider adding a comment to explain what the function does and what it returns. This can make it easier for others (or yourself!) to understand the code.
- * If you're working with large arrays, you might want to consider using a more efficient algorithm or data structure to store and process the values. However, for small arrays like the one in your example, the simple loop-based approach should be sufficient.

Overall, your code is simple and easy to understand, and the fix is a good one. With a few minor tweaks, it should be even better!

(a) Feedback with Llama3-8B and basic prompt.

Think about what you're actually trying to sum up in your loop. Are you trying to sum up the indices of the array, or the values themselves? Take a closer look at the condition in your if statement and see if it's doing what you intend it to do.

(b) Feedback with Llama3-8B and engineered prompt.

The problem with the code is where you are adding the sum. Rather than adding the value of the element, you are adding the index of that element.

(c) Feedback with Llama3-8B-fine-tuned and basic prompt.

Figure 1: Examples of feedback by Llama3-8B with different settings for a buggy program attempting the “Sum Positive Values” problem (see Figure 2). Feedback with the fine-tuned model is correct in contrast to the feedback given with the engineered prompt which contains redundant information that may confuse the student.

illustrated in Figure 1. This approach reduces reliance on manual prompt engineering while enhancing the human-like qualities of automated feedback in programming education. Our contributions are as follows:

- I. We introduce a fine-tuning method using learnersourced data to align AI-generated feedback with student-written feedback, reducing reliance on prompt engineering.
- II. We propose a flipped-role setup for feedback collection and define key attributes for analyzing feedback style.
- III. Our fine-tuned models produce feedback that is both more similar to student-written feedback and more accurate than basic and prompt-engineered approaches.
- IV. We publicly release our fine-tuning implementation to support future research in educational data mining.

2. RELATED WORK

Expert-designed feedback. Even prior to the advent of generative AI, significant efforts were made to provide automated feedback to students in introductory programming courses. Early work utilized unit test-based feedback in automated assessment systems, where students are shown specific error messages depending on which tests pass or fail [3, 11, 30]. Other studies relied on expert-crafted or learned rules to identify specific bugs, thereby providing tailored feedback [4, 15, 34, 40]. While these methods ensured high-quality feedback, they required substantial expert involvement and large training datasets, limiting scalability.

Crowdsourced programming feedback. To overcome the limitations of expert-designed feedback, researchers explored

crowdsourcing to collect errors and fixes contributed by the programming community [2, 9, 13, 24]. Initiatives like HelpMeOut [8] focused on collecting and matching code fixes for buggy programs, but these methods faced challenges in generalizability and struggled to adapt to new bug types. Learnersourcing builds on this concept by engaging students directly in creating feedback [36–38], which has been shown to generate valuable, relatable feedback and improve student learning [19, 35]. Learnersourcing remains relevant alongside generative AI [14], with recent studies exploring their synergy by examining student-written hints for programming bugs, both with and without AI assistance [39]. In contrast, our work focuses on leveraging learnersourced data to enhance generative AI models for better feedback generation.

Generative AI-powered programming feedback. The rise of generative models has opened new possibilities for automated feedback in education. Tools like GPT-3 and Codex [27] have been used to enhance compiler error messages and provide syntax bug fixes [20, 31]. Efforts to improve the quality of AI-generated feedback include incorporating symbolic information [12, 33, 42], validation mechanisms [31, 33], and retrieval-augmented generation (RAG) [7, 21]. However, these studies have primarily focused on generating feedback that meets the standards of expert educators, often overlooking student preferences. Our work bridges this gap by focusing on generating feedback that aligns with learner-sourced data, thereby enhancing the overall experience.

Language model fine-tuning. Recent work has explored fine-tuning generative models, particularly small open-access models like Llama3-8B [23] and Phi3-3.8B [1], to optimize for specific tasks in education [22, 43]. Fine-tuning with syn-

Problem Description

Sum Positive Values
 Define a function called SumPositiveValues() which is passed two inputs: an array of integers, and an integer indicating how many elements are in the array. The function should return the sum of all positive integers in the input array.

Buggy Program

```

1 int SumPositiveValues(int values[], int length) {
2     int sum = 0;
3     for (int i = 0; i < length; i++) {
4         if (values[i] > 0){
5             sum += i;
6         }
7     }
8     return sum;
9 }
    
```

(a) Scenario

Instruction for the Flipped Role Task

Imagine you are a tutor. A student in the course is asking for your help. Here is the problem description and the student's buggy C program. You should provide feedback so that the student can understand the issues in their buggy program and fix it. Provide your feedback in the textbox below.

Student-Written Feedback

The program needs to output the sum of all the positive values. This program is outputting the sum of all of the indexes of the values. To fix it make sure you add the values not the index i.

(b) Student acts in a flipped role as a tutor

Figure 2: (a) shows the scenario including the problem description and the buggy code. (b) shows the instruction given to the student asking them to act in the flipped role as a tutor and give feedback and a genuine example of student-written feedback.




Figure 3: Feedback-writing task (Sections 3.1 and 3.2)

thetically generated data has been shown to significantly improve feedback quality for programming tasks [16]. In contrast, we use real student-written hints for fine-tuning, aiming to mimic human-like styles while improving the accuracy and alignment of AI-generated feedback.

3. METHODOLOGY

This section details our methodology for collecting and analyzing learnersourced feedback, as well as generating and evaluating AI-generated feedback. First, we describe the flipped-role feedback-writing task used to collect student-written feedback. Next, we outline the dataset, including problems, buggy programs, and collected responses. We then present our approach for feedback generation using different prompting strategies and our fine-tuning techniques. Finally, we explain the evaluation setup, including the rubric and expert annotation process used to analyze the attributes of the learnersourced and AI-generated feedback.

3.1 Feedback-Writing Task

Flipped role task. To leverage learnersourcing, we designed a feedback-writing task where students take on the role of tutors, providing feedback on buggy code. Figure 3 illustrates the overall process, detailing the key components of the task. Students receive a problem description, a buggy implementation, and an instruction asking them to provide feedback from a tutor's perspective. This setup encourages students to analyze errors, articulate their reasoning, and mimic the peer feedback process in programming education. Figure 2 presents an example of this task, showing a specific problem along with the buggy code, the instruction given to the student, and a genuine example of student-written feedback. Since students typically have less programming

experience than real tutors, providing meaningful feedback can be challenging. To support them, we incorporated a structured pre-feedback activity to help them better understand the bug before writing feedback.

Pre-feedback activity. When presented with the problem and the buggy code, students first complete an activity where they are required to provide an input for a failing test case that would trigger a bug in the program, the buggy output on this input, and the correct output for this input. Based on whether they successfully complete this step, we categorize their understanding into two groups: *Understanding=1* (successful) and *Understanding=0* (unsuccessful). Regardless of their performance, all students are then shown a fixed version of the buggy code to help them better understand the problem and the bug. Finally, they write feedback on the buggy program from a tutor's perspective, aiming to explain the issue and suggest how to fix it. This structured process ensures students reflect on errors before providing guidance.

3.2 Student-Written Feedback Data

Set of feedback-writing tasks. To prepare tasks, we started with 3 programming problems and 10 buggy programs for each. The problems and buggy programs were pre-selected by experts to cover different introductory programming concepts and capture diverse types of bugs. For each task, the student is given a problem and a buggy program and asked to act in the flipped role of a tutor to write feedback for that buggy program. In total, we obtained 30 feedback-writing tasks. Below are the three problem descriptions:

- Problem 1: "Sum Positive Values" – Return the sum of the positive values in the input array.
- Problem 2: "Print Summary" – Print "Positive", "Negative", or "Equal", depending on whether there are more positive or negative numbers in the input array until the terminating value 0.
- Problem 3: "Print Average Rainfall" – Compute the average of non-negative integer values in the input array until the terminating value of -999.

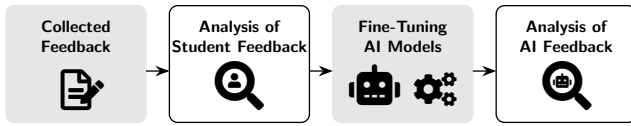


Figure 4: Data analysis workflow (Sections 3.3 and 3.4)

In terms of difficulty, Problem 1 is generally considered simpler, while Problems 2 and 3 are more challenging.

Course and students. The data collection process used a setup similar to a prior work [29]. Following this approach, we conducted our study in an introductory C programming course at the University of Auckland. There are about 750 students enrolled in the course, who typically have little to no prior experience in programming. To allow for a large number of students to work on the feedback-writing tasks, we developed a web application and deployed it as part of a laboratory exercise, conducted towards the end of the course. During this exercise, each student was given three feedback-writing tasks, one for each problem in a fixed order (Problem 1 → Problem 2 → Problem 3).

Data. Over 700 students completed all three tasks. Regarding providing a correct failing test case that identified a bug, 503 (67.9%) students succeeded in Problem 1, 308 (41.6%) in Problem 2, and 314 (42.4%) in Problem 3. After excluding empty entries, we collected 1920 feedback instances across all tasks. The feedback ranged from concise single-word responses to detailed explanations of up to 295 words. In general, the feedback consisted of complete well-structured sentences, as shown in Figure 2b (bottom).

3.3 AI-Generated Feedback Data

Technique for AI-generated feedback. To investigate AI-generated feedback, we use a technique grounded in literature, which involves leveraging a problem description and a buggy code, along with additional symbolic information (failing test case and a fixed version of the code) to generate feedback [32, 33, 42].

The fixed version of the code is typically obtained using a generative model with a separate prompt before asking the model for feedback.¹ Beyond this information, another important aspect is the instruction given to guide the model regarding the style of feedback it should generate. Here, we will investigate two prompting strategies:

- **Basic prompt:** This prompt uses a basic instruction for the model to generate feedback without any style guidance (see Figure 5a).
- **Engineered prompt:** This prompt adopts instructions from techniques shown to be effective, developed through expert engineering in existing work [33, 42] (see Figure 5b).

Generative models. Our analysis involves models from OpenAI’s GPT family, specifically GPT-3.5 Turbo [25, 27] and

¹As the focus of this study is to investigate AI-generated feedback in terms of style and structure, we used the same fixed code for all models when asking for feedback – this allows us to compare more directly the impact of prompt instructions and fine-tuning.

Basic Prompt for Feedback Generation

I am working on a C programming problem. The current program below is not working well. Can you help by giving feedback?

Problem description: {problem_description}
 Failing test cases: {failing_test_case}
 Buggy program: {buggy_program}
 Fixed program for the buggy program: {fixed_program}

Can you help by giving feedback?

(a) Basic prompt

Engineered Prompt for Feedback Generation

I am working on a C programming problem. The current program below is not working well. Can you help by giving feedback?

Problem description: {problem_description}
 Failing test cases: {failing_test_case}
 Buggy program: {buggy_program}
 Fixed program for the buggy program: {fixed_program}

1. Describe the bugs and provide an explanation along with fixes.
 2. Provide a concise hint about one bug in the buggy code. Do not give out the answer or any code. If there’s an obvious bug, direct to the location of the bug. If there’s a conceptual misunderstanding, offer a conceptual refresher. Limit your response for the hint to a sentence or two at most. Be as Socratic as possible, and be super friendly.

(b) Engineered prompt

Figure 5: Comparison of prompts using two kinds of instructions (highlighted in yellow) for generating feedback.

GPT-4 Turbo [26, 27], which have previously demonstrated their efficacy in providing feedback for debugging purposes [32, 33]. Additionally, we used small open-access models for fine-tuning, namely Llama3-8B [23] and Phi3-3.8B [1] due to their growing popularity in educational settings for a lower cost, convenience, and better data privacy [16, 17].

Fine-tuning setup. To examine the effects of model fine-tuning on generated feedback, we conducted supervised fine-tuning on Llama3-8B and Phi3-3.8B with student-written feedback². From the total of 1920 collected instances, we applied a filtering step to keep only those between 5 and 200 words, resulting in 1903 feedback instances. Our fine-tuning prompt mirrors the basic prompt as shown in Figure 5a but includes formatting instructions (i.e., asking the model to include in its answer a start and an end token). We picked hyperparameters based on existing literature [44].

3.4 Evaluation Setup

Evaluation rubric. Our aim is to first understand student-written feedback characteristics and then align generative models with them. To achieve this, we develop a detailed rubric capturing various attributes of feedback:

- **Correct** (binary): Indicates whether the feedback provides correct information that can help with debugging the code.

²The code used for fine-tuning these models is available at <https://github.com/machine-teaching-group/edm2025-humanizing-feedback>.

Table 1: Results for a sample of 300 instances of student-written feedback. We first break down the numbers per problem, then by the student’s ability to provide a test case (i.e., Understanding), and finally we combine the two conditions.

Problem{1/2/3}: Understanding{Any/1/0}	Sample Size	Correct %	Num. Words	Num. Sentences	Gives Fix %	Mentions Variables %	Mentions Lines %
Three problems:Understanding=Any	300	77.3	46.1	2.7	46.0	36.3	11.3
Problem 1:Understanding=Any	100	78.0	49.4	2.9	48.0	41.0	14.0
Problem 2:Understanding=Any	100	79.0	47.4	2.8	48.0	31.0	9.0
Problem 3:Understanding=Any	100	75.0	42.3	2.6	42.0	37.0	11.0
Three problems:Understanding=1	172	84.3	49.5	2.9	49.4	40.7	14.5
Three problems:Understanding=0	128	68.0	41.5	2.5	41.4	30.5	7.0
Problem 1:Understanding=1	74	85.1	50.3	2.9	51.4	40.5	14.9
Problem 1:Understanding=0	26	57.7	43.8	2.7	38.5	42.3	11.5
Problem 2:Understanding=1	47	83.0	57.2	3.4	53.2	36.2	10.6
Problem 2:Understanding=0	53	75.5	38.8	2.3	43.4	26.4	7.5
Problem 3:Understanding=1	51	84.3	41.5	2.4	43.1	45.1	17.6
Problem 3:Understanding=0	49	65.3	43.2	2.7	40.8	28.6	4.1

- *Num. words* (integer): Counts the number of whitespace-separated words in the feedback.
- *Num. sentences* (integer): Counts the number of sentences in the feedback, identified by end-of-sentence punctuation or new lines.
- *Gives fix* (binary): Indicates whether the feedback explicitly suggests how to fix the buggy program, either conceptually or by specifying changes.
- *Mentions variables* (binary): Indicates whether the feedback mentions specific variables from the buggy code.
- *Mentions lines* (binary): Indicates whether the feedback mentions line numbers in the code.

The attributes *Num. words* and *Num. sentences* were computed automatically. The remaining attributes were manually annotated, as discussed below.

Expert annotation for student-written feedback. We asked two experts to evaluate the feedback in a scheme similar to prior work [10, 31] according to the rubric described above. The experts assessed a random sample of 300 student-written feedback instances, with 100 instances for each of the three problems. First, each expert independently annotated 10 instances per problem to assess inter-rater reliability. They achieved a *substantial agreement* with a Cohen’s kappa of 0.63 [5].³ Afterward, one expert annotated the entire sample of student feedback.

Expert annotation for AI-generated feedback. As described in Section 3.2, there is a set of 10 buggy programs for each of the three problems. We generated the feedback using all of the models for each of the programs. The same two experts mentioned above independently annotated each instance based on the described rubric.

4. RESULTS

In this section, we analyze the characteristics of student-written feedback and evaluate the effectiveness of fine-tuning

³The evaluation rubric can be refined iteratively to improve the inter-rater agreement. Future work could explore rubric refinement to further improve reliability and robustness.

generative models to align with it. First, we examine the correctness, conciseness, and structure of student feedback, considering factors such as problem complexity and student understanding of the bug. Then, we compare AI-generated feedback across different prompting strategies and fine-tuned models, assessing how well they replicate student feedback in both style and accuracy.

4.1 Analysis of Student-Written Feedback

To understand the characteristics of student-written feedback, we analyze three key aspects: (a) whether students provide correct feedback and if they prefer concise hints or detailed explanations, (b) how problem complexity influences feedback characteristics, and (c) how a student’s understanding of the bug affects the quality and detail of their feedback.

Student feedback is mostly correct and concise. First, we explore the student-written feedback in terms of correctness and style. As shown by the first row in Table 1, the overall feedback correctness is high, approaching 80%. This indicates that students are generally capable of providing accurate feedback for programming bugs. We note that this high rate could be partly affected by the fact that we tried to increase feedback correctness by showing the students the fixed program, as described in Section 3. A feedback instance on average comprises only 46 words spanning fewer than 3 sentences, with slightly fewer than half of them explicitly mentioning how to fix the buggy code. This suggests that students generally favor short, concise feedback but some students are flexible and include actionable ideas for fixing the code. Additionally, we observe that students prefer using variable names rather than line numbers for localizing bugs and fixes.

Problem complexity has little effect on feedback style. Next, we investigate whether the complexity of the problem affects the type of feedback students provide. Despite the assumption that Problem 1 is simpler and Problems 2 and 3 are more complex, no clear differences emerged between the three problems in terms of feedback correctness as well as the percentages of giving fixes and mentioning variables or line numbers. The only noticeable variation is in the length of

Table 2: Results for the AI-generated feedback aggregated over two experts, compared to student-written feedback. We assessed the models for 10 buggy programs for each of the 3 problems. Rows are grouped by model family.

Feedback Source	Prompting Strategy	Num. Problems	Num. Programs per Problem	Correct %	Num. Words	Num. Sentences	Gives Fix %	Mentions Variables %	Mentions Lines %
Human students		3	100	77.3	46.1	2.7	46.0	36.3	11.3
GPT-4 Turbo	basic	3	10	81.7	411.3	47.7	96.7	100.0	1.7
GPT-4 Turbo	engineered	3	10	96.7	35.6	2.4	6.7	23.3	0.0
GPT-3.5 Turbo	basic	3	10	63.3	140.2	11.8	80.0	71.7	3.3
GPT-3.5 Turbo	engineered	3	10	90.0	20.1	1.4	0.0	15.0	1.7
Llama3-8B	basic	3	10	56.7	256.3	32.2	100.0	90.0	3.3
Llama3-8B	engineered	3	10	71.7	27.8	3.3	3.3	20.0	0.0
Phi3-3.8B	basic	3	10	55.0	210.6	35.1	100.0	58.3	0.0
Phi3-3.8B	engineered	3	10	80.0	25.7	1.9	5.0	5.0	0.0
Llama3-8B-fine-tuned	basic	3	10	86.7	47.7	2.7	71.7	40.0	1.7
Phi3-3.8B-fine-tuned	basic	3	10	88.3	68.0	3.9	98.3	60.0	1.7

the feedback, with the number of words and sentences tending to gradually decrease. We note that this trend might be attributed to the fixed order in which the problems were presented to the students.

Better understanding yields more targeted feedback. Finally, we analyze whether the students’ performance in understanding bugs affects their feedback style. The last eight rows in Table 1 illustrate the results concerning students’ understanding of the problem and buggy code, as indicated by their ability to provide a failing test case. Correctness is higher among students who provided a good test case for the bug, even though an example of a good test case and fixed code were provided afterward before they wrote feedback. Additionally, these students tend to offer longer and more detailed explanations, include fixes more frequently, and provide more targeted feedback by referencing specific aspects of the buggy code.

4.2 Fine-Tuning for Feedback Generation

We now analyze the impact of prompt engineering and fine-tuning on AI-generated feedback. First, we examine how prompting strategies influence the style and correctness of model-generated feedback. Then, we evaluate whether fine-tuning models on learnersourced data improves accuracy and alignment with student-written feedback.

Engineered prompts reduce verbosity. We first investigate the effects of prompting strategies on generated feedback, and how well the engineered prompt guides the model’s output toward resembling student feedback. The first nine rows in Table 2 present a comparison between feedback from students and various generative models, differentiating on base models and prompting strategies. It is evident that all base models using the basic prompt tend to be verbose and usually provide fixes. However, the engineered prompt substantially reduces verbosity and enhances correctness, more closely aligning the feedback style with that of students. Figures 1a and 1b show contrastive examples of the verbosity between the basic and engineered prompts. Because the engineered prompt is based on existing expert-crafted designs, it explicitly instructs AI models not to provide fixes. While this enables direct comparison with prior work, it may not always align with how students naturally write feedback.

Fine-tuning improves alignment and accuracy. Next, we answer whether fine-tuning open-access models with student-written feedback can potentially replace more complex prompt engineering approaches for generating feedback. We investigated the effectiveness of fine-tuning Llama3-8B and Phi3-3.8B. The last two rows in Table 2 indicate that both models produced short feedback, more closely resembling the style of students. The feedback generated by Llama3-8B-fine-tuned is closest to student-written feedback in terms of length and frequency of mentioning variables. Another observation is that these AI models tend to use line numbers less frequently than students, preferring to name variables instead. Remarkably, the correctness of AI-generated feedback improves after fine-tuning, surpassing that of the base model with the engineered prompt even though some of the feedback instances used for fine-tuning are incorrect. This improvement may result from the models becoming more familiar with the domain and generating shorter output, reducing the chances of errors.

5. CONCLUDING DISCUSSIONS

This study explored how learnersourced data can enhance AI-generated feedback, aligning it more closely with student communication styles. We first characterized student-written feedback on buggy code, focusing on correctness, length, inclusion of explicit fixes, and methods of bug identification. Our findings reveal that students tend to write concise feedback while remaining flexible in suggesting fixes, highlighting key patterns that AI-generated feedback should replicate. Building on these insights, we fine-tuned generative models on learnersourced data, reducing reliance on expert prompt engineering. This approach improves the adaptability of AI-generated feedback, making it more peer-like and responsive to diverse student needs. Notably, even when some student feedback was incorrect, fine-tuning still led to substantial improvements in accuracy, showcasing the potential of student-centered fine-tuning for scalable and personalized feedback systems in programming education.

While our findings are promising, several limitations suggest avenues for future research. First, we did not evaluate the models’ effectiveness in real classroom settings. Deploying these models in real classrooms and gathering student reflections on their utility and relevance would provide insights

into their practical impact. Second, students in our study were given a test case and a corrected program to maximize feedback accuracy. While this ensured high correctness rates (close to 80%), it may not reflect real-world scenarios where such guidance is unavailable. Future work could examine how feedback quality and correctness change when students must identify bugs without prior guidance, evaluating model robustness in more realistic settings. Third, we did not explore the potential of fine-tuning models for specific courses, which could improve feedback relevance and better support course-specific learning needs. Employing learnersourced data in this context also introduces ethical considerations, such as ensuring data protection, respecting intellectual property rights, and transparently communicating how content created by students is utilized. Although our analysis did not reveal that models inherited misconceptions from student-written feedback, future applications should explicitly monitor and mitigate this risk. Addressing these ethical and methodological aspects thoughtfully can strengthen the reliability and effectiveness of fine-tuned generative models. Ultimately, combining crowdsourcing with fine-tuning remains a promising and scalable approach to delivering personalized, high-quality feedback, especially valuable in low-resource educational settings.

Acknowledgments

Juho Leinonen acknowledges funding by Research Council of Finland (Academy Research Fellow grant number 356114). Michael Liut acknowledges funding by NSERC Discovery Grant #RGPIN-2024-04348. Funded/Co-funded by the European Union (ERC, TOPS, 101039090). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] M. I. Abdin et al. Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone. *CoRR*, abs/2404.14219, 2024.
- [2] A. Al-batlai, M. Abdullah-Al-Wadud, and M. A. Hosain. A Review on Recommending Solutions for Bugs Using Crowdsourcing. In *Saudi Computer Society National Computer Conference (NCC)*, 2018.
- [3] K. M. Ala-Mutka. A Survey of Automated Assessment Approaches for Programming Assignments. *Computer Science Education*, 2005.
- [4] B. A. Becker. An Effective Approach to Enhancing Compiler Error Messages. In *Proc. of the Technical Symp. on Computer Science Education (SIGCSE)*, 2016.
- [5] J. Cohen. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 1960.
- [6] P. Denny, S. Gulwani, N. T. Heffernan, T. Käser, S. Moore, A. N. Rafferty, and A. Singla. Generative AI for Education (GAIED): Advances, Opportunities, and Challenges. *CoRR*, abs/2402.01580, 2024.
- [7] M. Frazier, K. Damevski, and L. Pollock. Customizing ChatGPT to Help Computer Science Principles Students Learn Through Conversation. In *Proceedings of the Conference on Innovation and Technology in Computer Science Education (ITiCSE)*, 2024.
- [8] B. Hartmann, D. MacDougall, J. Brandt, and S. R. Klemmer. What Would Other Programmers Do: Suggesting Solutions to Error Messages. In *Proceedings of the International Conference on Human Factors in Computing Systems (CHI)*, 2010.
- [9] A. Head, E. L. Glassman, G. Soares, R. Suzuki, L. Figueredo, L. D’Antoni, and B. Hartmann. Writing Reusable Code Feedback at Scale with Mixed-Initiative Program Synthesis. In *Proceedings of the Conference on Learning @ Scale (L@S)*, 2017.
- [10] A. Hellas, J. Leinonen, S. Sarsa, C. Koutcheme, L. Kujanpää, and J. Sorva. Exploring the Responses of Large Language Models to Beginner Programmers’ Help Requests. In *Proceedings of the Conference on International Computing Education Research (ICER)*, 2023.
- [11] P. Ihantola, T. Ahoniemi, V. Karavirta, and O. Sepälä. Review of Recent Systems for Automatic Assessment of Programming Assignments. In *Proceedings of the Koli Calling International Conference on Computing Education Research*, 2010.
- [12] M. A. Islam et al. MapCoder: Multi-Agent Code Generation for Competitive Problem Solving. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [13] H. Jiang et al. Toward Better Summarizing Bug Reports with Crowdsourcing Elicited Attributes. *IEEE Transactions on Reliability*, 2018.
- [14] H. Khosravi, P. Denny, S. Moore, and J. Stamper. Learnersourcing in the Age of AI: Student, Educator and Machine Partnerships for Content Creation. *Computers and Education: Artificial Intelligence*, 2023.
- [15] T. Kohn and B. Manaris. Tell Me What’s Wrong: A Python IDE with Error Messages. In *Proceedings of the Technical Symposium on Computer Science Education (SIGCSE)*, 2020.
- [16] N. Kotalwar, A. Gotovos, and A. Singla. Hints-In-Browser: Benchmarking Language models for Programming Feedback Generation. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*, 2024.
- [17] C. Koutcheme, N. Dainese, S. Sarsa, A. Hellas, J. Leinonen, and P. Denny. Open source language models can provide feedback: Evaluating llms’ ability to help students using gpt-4-as-a-judge. In *Proceedings of the Conference on Innovation and Technology in Computer Science Education (ITiCSE)*, 2024.
- [18] H. Kumar, I. Musabirov, M. Reza, J. Shi, A. Kuzminykh, J. J. Williams, and M. Liut. Impact of Guidance and Interaction Strategies for LLM Use on Learner Performance and Perception. *CoRR*, abs/2310.13712, 2023.

- [19] H. Lahza et al. Analytics of Learning Tactics and Strategies in an Online Learnersourcing Environment. *Journal of Computer Assisted Learning*, 2023.
- [20] J. Leinonen, A. Hellas, S. Sarsa, B. N. Reeves, P. Denny, J. Prather, and B. A. Becker. Using Large Language Models to Enhance Programming Error Messages. In *Proceedings of the Technical Symposium on Computer Science Education (SIGCSE)*, 2023.
- [21] R. Liu, C. Zenke, C. Liu, A. Holmes, P. Thornton, and D. J. Malan. Teaching CS50 with AI: Leveraging Generative Artificial Intelligence in Computer Science Education. In *Proceedings of the Technical Symposium on Computer Science Education (SIGCSE)*, 2024.
- [22] S. Liu, Z. Yu, F. Huang, Y. Bulbulia, A. Bergen, and M. Liut. Can Small Language Models With Retrieval-Augmented Generation Replace Large Language Models When Learning Computer Science? In *Proceedings of Innovation and Technology in Computer Science Education (ITiCSE)*, 2024.
- [23] Meta. Llama-3. <https://ai.meta.com/blog/meta-llama-3/>, 2024.
- [24] D. Mujumdar, M. Kallenbach, B. Liu, and B. Hartmann. Crowdsourcing Suggestions to Programming Problems for Dynamic Web Development Languages. In *Extended Abstracts on Human Factors in Computing Systems (CHI)*, 2011.
- [25] OpenAI. ChatGPT. <https://openai.com/blog/chatgpt>, 2023.
- [26] OpenAI. GPT-4 Technical Report. *CoRR*, abs/2303.08774, 2023.
- [27] OpenAI. OpenAI Platform Models. <https://platform.openai.com/docs/models>, 2024.
- [28] B. Otaki and O. Lindwall. Generative AI and the Human Touch: Investigating the Changing Landscape of Feedback in Higher Education. In *Proceedings of the International Conference of the Learning Sciences (ICLS)*, 2024.
- [29] V. Padurean, P. Denny, and A. Singla. BugSpotter: Automated Generation of Code Debugging Exercises. In *Proceedings of the Technical Symposium on Computer Science Education (SIGCSE)*, 2025.
- [30] J. Paiva et al. Automated Assessment in Computer Science Education: A State-of-the-Art Review. *Transactions on Computing Education (TOCE)*, 2022.
- [31] T. Phung, J. Cambronero, S. Gulwani, T. Kohn, R. Majumdar, A. Singla, and G. Soares. Generating High-Precision Feedback for Programming Syntax Errors using Large Language Models. In *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2023.
- [32] T. Phung, V. Pădurean, J. Cambronero, S. Gulwani, T. Kohn, R. Majumdar, A. Singla, and G. Soares. Generative AI for Programming Education: Benchmarking Chatgpt, GPT-4, and Human Tutors. In *Proceedings of the Conference on International Computing Education Research (ICER) - Volume 2*, 2023.
- [33] T. Phung, V. Padurean, A. Singh, C. Brooks, J. Cambronero, S. Gulwani, A. Singla, and G. Soares. Automating Human Tutor-Style Programming Feedback: Leveraging GPT-4 Tutor Model for Hint Generation and GPT-3.5 Student Model for Hint Validation. In *Proceedings of the International Learning Analytics and Knowledge Conference (LAK)*, 2024.
- [34] C. Piech, J. Huang, A. Nguyen, M. Phulsuksombati, M. Sahami, and L. J. Guibas. Learning Program Embeddings to Propagate Feedback on Student Code. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- [35] N. Pirttinen, P. Denny, A. Hellas, and J. Leinonen. Lessons Learned from Four Computing Education Crowdsourcing Systems. *IEEE Access*, 2023.
- [36] N. Pirttinen and J. Leinonen. Can Students Review Their Peers? Comparison of Peer and Instructor Reviews. In *Proceedings of the Annual Conference on Innovation and Technology in Computer Science Education (ITiCSE)*, 2022.
- [37] A. Singh, C. Brooks, and S. Doroudi. Learnersourcing in Theory and Practice: Synthesizing the Literature and Charting the Future. In *Proceedings of the Conference on Learning @ Scale (L@S)*, 2022.
- [38] A. Singh, C. Brooks, Y. Lin, and W. Li. What’s in It for the Learners? Evidence from a Randomized Field Experiment on Learnersourcing Questions in a MOOC. In *Proceedings of the Conference on Learning @ Scale (L@S)*, 2021.
- [39] A. Singh, C. Brooks, X. Wang, W. Li, J. Kim, and D. Wilson. Bridging Learnersourcing and AI: Exploring the Dynamics of Student-AI Collaborative Feedback Generation. In *Proceedings of the International Learning Analytics and Knowledge Conference (LAK)*, 2024.
- [40] R. Singh et al. Automated Feedback Generation for Introductory Programming Assignments. In *Proceedings of the Conference on Programming Language Design and Implementation (PLDI)*, 2013.
- [41] S. Wang, J. C. Mitchell, and C. Piech. A Large Scale RCT on Effective Error Messages in CS1. In *Proceedings of the Technical Symposium on Computer Science Education (SIGCSE)*, 2024.
- [42] J. Zamfirescu-Pereira, L. Qi, B. Hartmann, J. DeNero, and N. Norouzi. Conversational Programming with LLM-Powered Interactive Support in an Introductory Computer Science Course. *NeurIPS’23 Workshop on Generative AI for Education (GAIED)*, 2023.
- [43] Y. Zheng, R. Zhang, J. Zhang, Y. Ye, Z. Luo, and Y. Ma. LlamaFactory: Unified Efficient Fine-Tuning of 100+ Language Models. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [44] C. Zhou et al. LIMA: Less Is More for Alignment. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2023.