McMining: Automated Discovery of Misconceptions in Student Code

Anonymous ACL submission

Abstract

When learning to code, students often develop misconceptions about various programming language concepts. These can not only lead to bugs or inefficient code, but also slow down the learning of related concepts. In this paper, we introduce McMining, the task of mining programming misconceptions from samples of code from a student. To enable the training and evaluation of McMining systems, we develop an extensible benchmark dataset of misconceptions together with a large set of code samples where these misconceptions are manifested. We then introduce two LLM-based McMiner approaches and through extensive evaluations show that models from the Gemini, Claude, and GPT families are effective at discovering misconceptions in student code.
https://anonymous.4open.science/r/mcminer-843E

1 Introduction and Motivation

A misconception is a belief in a false statement, in short a false belief, such as believing that the Earth is flat or that real numbers are countable. In the CS domain, programming misconceptions are beliefs about programming concepts that are not warranted by the programming language definition. We include here beliefs about the syntax and semantics of programming language constructs and builtin functions included with the language. For example, a common misconception in Python is that the `range(n)` function produces integers starting at 1 and ending at n, which may potentially be the cause for the bug shown in Figure 1.

Programming misconceptions are often the cause of bugs, which is obviously detrimental. Furthermore, in an educational context, student learning is impeded significantly when misconceptions slow down their ability to solve problems correctly, e.g. writing code that passes test cases. Misconceptions, knowledge gaps, and other types of difficulties (Qian and Lehman, 2017) encountered by students make it difficult to learn one concept correctly, which then makes it harder to acquire other closely linked concepts, propagating and compounding throughout a course in a snowballing manner (Robins, 2010). To maintain positive learning momentum, it is therefore essential that misconceptions are identified and fixed as early as possible. Typically, identifying a misconception is done by the students themselves, alone or ideally under Socratic guidance through conversations with an instructor or teaching assistant (Al-Hossami et al., 2024). However, this process can be cognitively demanding, and, due to insufficient TA resources (Yadav et al., 2016), not fast enough to keep up with the volume of concepts introduced in a course.

➢ **Problem description**:

Write the `factorial(n)` function that computes the factorial `n!` defined as:

```
0! = 1
n! = n x (n - 1)!
```

If the input `n` is negative, the function should return `0`.

➢ **Student code**

```
def factorial (n):
  if n < 0:
    return 0
  fact = 1
  for i in range(n):
    fact = fact * i
  return fact
```

➢ **Potential misconception**:

`range(n)` produces values from 1 to `n` inclusive.

Figure 1: Problem-code pair that exhibits a *potential* programming misconception about the `range` function.

Recognizing the importance of identifying misconceptions early, in this paper we introduce McMining, the novel task of mining programming misconceptions from code samples produced by a student over time. Correspondingly, we describe the development of a benchmark dataset of programming misconceptions together with a large set of problem descriptions and corresponding code samples that exhibit these misconceptions (Section 3). We then introduce two versions of

an LLM-based tool for mining misconceptions, MCMINER-S that identifies potential misconceptions in one code sample at a time, and MCMINER-M which attempts to identify misconceptions as patterns exhibited by multiple code samples from a student (Section 4). We present positive results of the two version of the MCMINER tool on the benchmark dataset, using different LLMs (Section 5). The paper ends with conclusions and thoughts on future work (Section 6).

2 Related Work

Traditional automated systems (Johnson and Soloway, 1984; Sirkia and Sorva, 2012; Evans et al., 2023) rely on hand-crafted rules targeting a limited set of specific misconceptions. Similarly, modern LLM-based approaches such as (Lee et al., 2024) focus on classification of logical errors into a predefined set of categories. Overall, targeting only predefined or common misconceptions is a significant limitation, considering that educator beliefs about common errors can diverge significantly from actual student patterns (Brown and Altadmri, 2014). In parallel, other approaches aim to identify code segments that contain logical errors, such as (Mens et al., 2021; Hoq et al., 2025), *inter alia*. While knowing the location of logical errors is important, current approaches do not articulate what the actual misconception is about, which can be argued to be equally important from an educational standpoint. In this context, the MCMINING task is novel in that it requires not only the identification of potentially unknown misconceptions from consistent samples of student code, but also articulating their description in terms of a false belief about a programming concept.

3 Task Definition and Benchmark Dataset

Given a set of problems and the corresponding code samples produced by the student, misconception mining is the task of identifying any potential programming misconception that is exhibited in their code. More formally, the input to the misconception mining task consists of a set PC of problem-code pairs (p,c):

$$PC=\{(p_1,c_1),(p_2,c_2),\ldots,(p_N,c_N)\}=\{(p_n,c_n)\}_1^N \quad (1)$$

where each pair (p_n, c_n) contains a problem description p_n and the corresponding coding solution c_n provided by the student. Figure 1 shows an example (p, c) pair, where the code illustrates a common misconception about the range function. When presented with this and other (p,c) pairs that exhibit the same misconception, the model is expected to generate a natural language description of the misconception. Note that all programming misconceptions identified by the model are considered to be *potential* misconceptions. Looking at the example in Figure 1, it is possible that the student has the correct knowledge about the `range(n)` function and that the bug in their code is caused by mistakenly using `i` instead of `(i + 1)` on line 6. The more samples of buggy code that utilize the `range` function are in the input set PC, the more likely it is that the student who wrote that code holds this misconception.

While misconceptions are known to often cause bugs, they do not necessarily do so. For example, a student may incorrectly believe that all local variables need to have a one letter name. Correspondingly, we classify misconceptions into two disjoint categories, *harmful* vs. *benign*, depending on whether they cause or not a bug in the code.

To enable the evaluation of misconception mining tools, we developed a benchmark dataset $\mathcal{D}$ of K student submission sets, where each student submission set $PC^{(k)}$ contains problem-code pairs created such that a subset of the pairs exhibit a certain programming misconception $m^{(k)}$. More formally, the dataset contains the following components:

$$\mathcal{D}=\left\{\left(PC^{(k)}, mc^{(k)}\right)\right\}_1^K; \quad PC^{(k)}=\left\{\left(p_n^{(k)}, c_n^{(k)}\right)\right\}_1^{N_k} \quad (2)$$

Each example in $\mathcal{D}$ consists of two parts: the set of (p,c) samples $PC^{(k)}$ as the observed *input*, and the misconception $mc^{(k)}$ as the target *label*.

To develop $\mathcal{D}$, we first created a set $\mathcal{MC}=\left\{mc^{(m)}\right\}_1^M$ of $M=67$ misconceptions, of which 64 are common, documented misconceptions, and 3 are artificial misconceptions that are included to enable evaluation on novel misconceptions. We then created a set $\mathcal{PS}=\{(p_n,S_n)\}_1^N$ of $N=501$ problem-solution pairs (p,S), where for each problem p, the set S contains one or more *correct* coding solutions for that problem.

Given $\mathcal{MC}$ and $\mathcal{PS}$, the benchmark dataset $\mathcal{D}$ is created as shown in Algorithm 1. An essential component of this procedure is shown at step 8, where a misconception is injected into the correct coding solution of a problem in order to create code that exhibits that misconception. In Section 3.1 below we describe the MCINJECT tool that we developed for this purpose.

Algorithm 1 Benchmark Dataset Generation

Initialize dataset $\mathcal{D} \leftarrow \emptyset$
for each student $k \in \{1,2,...,K\}$ **do**
  Select misconception $mc^{(k)} \in \mathcal{MC}$
  Initialize problem-code set $PC^{(k)} \leftarrow \emptyset$
  Sample from $\mathcal{PS}$ a set of N_k problem-solution pairs $PS^{(k)} = \left\{ \left(p_n^{(k)}, S_n^{(k)} \right) \right\}_1^{N_K}$
  for each $\left(p_n^{(k)}, S_n^{(k)} \right) \in PS^{(k)}$ **do**
    Sample a correct solution $s_n^{(k)} \in S_n^{(k)}$
    MCINJECT $mc^{(k)}$ in $s_n^{(k)}$ to obtain code $c_n^{(k)}$
    Add problem-code pair $\left(p_n^{(k)}, c_n^{(k)} \right)$ to $PC^{(k)}$
  Add $\left(PC^{(k)}, mc^{(k)} \right)$ to dataset $\mathcal{D}$
return dataset $\mathcal{D}$ of K misconception-labeled code sets

Misconceptions	67
Problem-Solution pairs used	25
Code samples	**1,675**
Exhibiting misconceptions	1,063
Showing no misconception	612
Bags of code samples	**339**
Bags with misconceptions	279
Bags with correct code only	60
Code samples per bag	4–8

Table 1: Benchmark dataset statistics.

Table 1 summarizes the key statistics of our benchmark dataset. We organized these samples into 339 code sets, or bags of code samples, each containing between $N_k = 4$ to $N_k = 8$ code samples. Importantly, 60 bags contain only correct code samples without any misconceptions, serving as negative examples to test the system's ability to correctly identify when no misconceptions are present.

3.1 The McInject Tool

The misconception injection tool MCINJECT takes as input the description of a computational problem p, a correct solution code s, and the description of a programming misconception mc. As output, it produces a code c as a version of s that is modified to exhibit the misconception mc. The code c should look as if written by a student holding misconception mc who tries to solve problem p. The tool is implemented using Claude Sonnet-4.5 with extended thinking, with a structured prompt and in-context learning examples. When a misconception cannot be meaningfully applied to a given solution code, MCINJECT is instructed to indicate so, rather than force inappropriate modifications. An LLM-as-judge evaluation shows that 90.3% of the code samples successfully exhibit their target misconceptions. Note that this is a conservative estimate of MCINJECT's performance, since the LLM-as-judge sometimes discards code samples that exhibit benign misconceptions, where the code is both correct and natural. Further manual evaluation of 88 code samples shows a 96.6% agreement between LLM judgment and human assessment, confirming the reliability of the LLM-as-judge. Appendix E provides further detailed descriptions of the tool development, including the prompt, LLM hyper-parameters, and evaluation setting.

Algorithm 2 MCMINER-S

Input: Bag of pairs $PC = \{(p_1, c_1), (p_2, c_2), \ldots, (p_N, c_N)\}$
Output: Potential misconception $\hat{mc}$ for the bag

1: Initialize set of misconceptions $M \leftarrow \emptyset$
2: **for** each problem-code pair $(p_j, c_j) \in PC$ **do**
3: Identify potential misconception mc_j for (p_j, c_j)
4: If found, add mc_j to M
5: **for** each misconception $mc \in M$ **do**
6: Let $count(mc) = |\{(p_j, c_j) \in PC | mc_j = mc\}|$
7: **if** all counts are 0 **then** ▷ *no misconception found*
8: ***return*** ϵ
9: ***else*** ▷ *some misconception found*
10: ***return*** $\hat{mc} = \underset{mc \in M}{\operatorname{argmax}}\ count(mc)$

4 The McMiner Tools

As defined in Section 3, in the misconception mining task the input consists of a bag of one or more problem-code pairs from a student, and the output is a potential misconceptions exhibited in the code samples. To solve this task, we developed two variants of an LLM-based mining tool: a *single instance* variant and a *multiple instance* variant.

The single instance variant MCMINER-S proceeds in two stages, as shown in Algorithm 2. First, an LLM is instructed to identify potential misconceptions in each problem-code pair (steps 2 to 4), using the prompt shown in Appendix F. Then, out of all the identified misconceptions, the one found in the most problem-code pairs is returned (steps 5 to 10).

In the multiple instance variant MCMINER-M, an LLM is given the entire bag of program-code pairs and is instructed to identify the misconception that is shared by the largest number of code samples in the bag, using the prompt shown in Appendix F. Compared to MCMINER-S, MCMINER-M has the advantage of looking at multiple code samples that may exhibit the same potential misconception pattern, which in theory should enable more precision in identifying misconceptions.

	MCMINER-M				MCMINER-S
Language Model	**Precision**	**Recall**	**F1-Score**	**Accuracy**	**Accuracy**
OpenAI o3-mini (low-effort)	85.6%	67.5%	75.5%	75.5%	76.9%
OpenAI o3-mini (medium-effort)	83.3%	70.8%	76.5%	76.9%	**78.5%**
Anthropic Claude Sonnet-4.5	79.1%	77.7%	78.4%	78.7%	66.4%
Anthropic Claude Sonnet-4.5 + Reasoning	83.8%	77.2%	**80.3%**	**82.0%**	68.7%
Gemini 2.5-flash	79.7%	56.2%	65.9%	61.1%	74.0%
Gemini 2.5-flash + Reasoning	77.0%	76.7%	76.8%	76.9%	69.4%

Table 2: MCMINER results for multiple-instance mining (left) and single-instance mining (right).

5 Experimental Evaluations

We used the benchmark dataset to evaluate the effectiveness of the two McMiner tools when instantiated using three LLMs: OpenAI's o3-mini models, Anthropic's Claude-Sonnet-4.5 models, and Gemini 2.5-Flash models. The experimental setup for each model is detailed in Appendix G.

Table 2 present the performance of different LLMs in the multiple and single instance mining settings. For computing accuracy, precision, recall, F1-score, we define: *true positives* are cases where the Ground Truth (GT) misconception matches the prediction, or the prediction is a validated novel misconception; *true negatives* are cases where both GT and prediction contain no misconception; *false positives* are cases where the prediction does not match GT and fails validation; and *false negatives* are cases where GT contains a misconception but the model predicted none. Note that when a validated novel misconception is found but GT contains a different misconception, this counts as both a true positive (for the valid discovery) and as a false negative (for missing GT).

Multi-instance mining (MCMINER-M) substantially outperforms single-instance mining (MCMINER-S), with the best model (Claude Sonnet-4.5 + Reasoning) achieving 82.0% accuracy for MCMINER-M compared to 78.5% for the best MCMINER-S model (o3-mini medium-effort). This demonstrates the value of analyzing multiple code samples simultaneously to identify consistent patterns with higher confidence. Enabling reasoning capabilities consistently improves MCMINER-M performance across model families: Claude (78.7% to 82.0%), Gemini (61.1% to 76.9%), and o3-mini (75.5% to 76.9% with medium effort).

The best MCMINER-M model achieves 93.3% accuracy on bags with misconceptions but only 59.3% on correct-only bags, indicating higher false positive rates on correct code.

MCMINER identifies novel misconceptions effectively: Claude Sonnet-4.5 with reasoning discovered 41 novel true positives (12.1% of bags), including: "The student believes that the division operator `/` and integer division operator `//` are interchangeable or that `/` will automatically return an integer when the result is a whole number."

Misconception difficulty varies substantially: syntax errors (e.g., `=` vs. `==`) achieve >95% accuracy, while subtle benign misconceptions (e.g., unnecessary parentheses) and subtle out-of-distribution cases (48% for vowel-based naming) prove more challenging, specially for single-instance mining. Theory of mind misconceptions (e.g., operator precedence) achieve 42-50% accuracy. Multi-instance mining improves detection of difficult cases.

6 Conclusion

This paper introduces MCMINER, an LLM-based tool for automatically discovering programming misconceptions in student code. Our evaluation on a benchmark of 1,063 code samples exhibiting 67 misconceptions demonstrates strong performance, with the best multi-instance model achieving 82.0% accuracy. Critically, MCMINER can identify emerging and rare student misconceptions not captured in existing taxonomies, discovering 41 novel misconceptions in our evaluation. This capability opens the door to uncovering rare, student-specific misconceptions and enables more personalized learning experiences.

Future work should systematically evaluate MCMINER on real student submissions and investigate pedagogical integrations.

7 Limitations

Our benchmark assumes each code bag exhibits a single primary misconception, while real students often hold multiple interacting misconceptions

simultaneously. Additionally, the evaluation relies on high-quality synthetic data rather than authentic student submissions. The evaluation is limited to Python; generalization to other programming languages remains unexplored. Our misconception taxonomy contains 67 misconceptions, the largest collection to our knowledge, and is designed to be easily extended, but cannot capture all possible student misconceptions.

Improving performance is an avenue for future research, particularly, reducing the false positive rate on correct code and its struggle with detecting subtle misconceptions specially those that require code tracing and theory of mind represent important areas for future research to improve precision.

MCMINER can only detect misconceptions related to programming that manifest as consistent code patterns. Misconceptions that are not related to programming, such as a student's misinterpreting the problem requirements, or having misconceptions regarding programming concepts like compilers, operating systems, or hardware, are not detectable by MCMINER.

Finally, the LLM-based approach incurs computational costs that may limit large-scale deployment in resource-constrained educational environments.

Deployment of MCMINER in real classrooms can raise pedagogical and privacy considerations. Over-reliance on automated detection may reduce instructors' diagnostic skills and discourage valuable Socratic dialogue with students. The system risks premature intervention before students have the opportunity for self-correction through debugging, a critical part of learning. Additionally, processing student code through third-party LLM APIs raises data privacy concerns and requires careful consideration of educational data protection regulations such as FERPA.

References

Erfan Al-Hossami, Razvan Bunescu, Justin Smith, and Ryan Teehan. 2024. Can language models employ the socratic method? experiments with code debugging. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, SIGCSE 2024, page 53–59, New York, NY, USA. Association for Computing Machinery.

Erfan Al-Hossami, Razvan Bunescu, Ryan Teehan, Laurel Powell, Khyati Mahajan, and Mohsen Dorodchi. 2023. Socratic questioning of novice debuggers: A benchmark dataset and preliminary evaluations. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 709–726, Toronto, Canada. Association for Computational Linguistics.

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and 1 others. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.

Neil CC Brown and Amjad Altadmri. 2014. Investigating novice programming mistakes: educator beliefs vs student data. In *Proceedings of the 2014 conference on Innovation & technology in computer science education*, pages 43–50.

Luca Chiodini, Igor Moreno Santos, Andrea Gallidabino, Anya Tafliovich, André L. Santos, and Matthias Hauswirth. 2021. A curated inventory of programming language misconceptions. In *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education V. 1*, ITiCSE '21, page 380–386, New York, NY, USA. Association for Computing Machinery.

Adrian de Freitas, Joel Coffman, Michelle de Freitas, Justin Wilson, and Troy Weingart. 2023. Falcon-code: A multiyear dataset of python code samples from an introductory computer science course. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*, SIGCSE 2023, page 938–944, New York, NY, USA. Association for Computing Machinery.

Andrew Ettles, Andrew Luxton-Reilly, and Paul Denny. 2018. Common logic errors made by novice programmers. In *Proceedings of the 20th Australasian Computing Education Conference*, ACE '18, page 83–89, New York, NY, USA. Association for Computing Machinery.

Abigail Evans, Zihan Wang, Jieren Liu, and Mingming Zheng. 2023. Side-lib: A library for detecting symptoms of python programming misconceptions. In *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1*, ITiCSE 2023, page 159–165, New York, NY, USA. Association for Computing Machinery.

Muntasir Hoq, Ananya Rao, Reisha Jaishankar, Krish Piryani, Nithya Janapati, Jessica Vandenberg, Bradford Mott, Narges Norouzi, James Lester, and Bita Akram. 2025. Automated Identification of Logical Errors in Programs: Advancing Scalable Analysis of Student Misconceptions. pages 90–103.

W Lewis Johnson and Elliot Soloway. 1984. Proust: An automatic debugger for pascal programs. In *Proceedings of the National Conference on Artificial Intelligence*, pages 181–184.

Yanggyu Lee, Suchae Jeong, and Jihie Kim. 2024. Improving LLM Classification of Logical Errors by Integrating Error Relationship into Prompts. In *Generative Intelligence and Intelligent Tutoring Systems:*

20th International Conference, ITS 2024, Thessaloniki, Greece, June 10–13, 2024, Proceedings, Part I, pages 91–103, Berlin, Heidelberg. Springer-Verlag.

Kim Mens, Siegfried Nijssen, and Hoang-Son Pham. 2021. The good, the bad, and the ugly: Mining for patterns in student source code. In *Proceedings of the 3rd International Workshop on Education through Advanced Software Engineering and Artificial Intelligence (EASEAI '21)*, pages 1–8. ACM.

Yizhou Qian and James Lehman. 2017. Students' Misconceptions and Other Difficulties in Introductory Programming: A Literature Review. *ACM Trans. Comput. Educ.*, 18(1):1:1–1:24.

Anthony Robins. 2010. Learning edge momentum: a new account of outcomes in CS1. *Computer Science Education*, 20(1):37–71. Publisher: Routledge _eprint: https://doi.org/10.1080/08993401003612167.

Teemu Sirkia and Juha Sorva. 2012. Exploring programming misconceptions: An analysis of student mistakes in visual program simulation exercises. In *Proceedings of the 12th Koli Calling International Conference on Computing Education Research*, pages 19–28.

Aman Yadav, Sarah Gretter, Susanne Hambrusch, and Phil Sands. 2016. Expanding computer science education in schools: understanding teacher experiences and challenges. *Computer Science Education*, 26(4):235–254.

A McInject Qualitative Analysis

We conducted a qualitative analysis of MCINJECT's generated code samples to understand performance variation across misconception types.

A.1 Success Patterns

Several misconceptions achieved 100% exhibit rates. Analysis reveals these fall into two primary categories:

Syntax-based misconceptions produce clear, unambiguous code patterns that directly violate Python syntax rules. For example, (*Student believes that the = operator is used for equality comparison in conditional statements*) and (*Student believes that colons (:) are used for variable assignment*) create syntax errors that are straightforward to inject and verify. Similarly,(*Student believes that a function can be defined without the def keyword*) removes a required keyword, making the misconception immediately apparent.

Clear misuse misconceptions involve well-defined incorrect usage of language constructs where the distinction between correct and incorrect usage is unambiguous. (*Student believes that a print statement must be used to return a value from a function*) exemplifies this pattern—replacing `return` with `print` creates a distinct behavioral signature that is both easy to inject and verify. (*Student believes that functions are called using square brackets like list indexing*). This misconception similarly produces code with an unmistakable incorrect pattern.

A.2 Failure Patterns

Conversely, misconceptions with exhibit rates below 70% reveal two distinct failure modes that highlight the challenges in generating realistic student-like code.

A.2.1 Ambiguous Correctness

Some misconceptions can manifest in code that is both correct and natural, making it difficult to determine whether the code reflects the misconception or simply represents a valid solution approach. Consider the misconception (*Student believes that loop iteration requires manual counter tracking with an explicit variable to access element indices*) for the problem *"Write a function to find the next smallest palindrome greater than a given number"* MCINJECT generated:

```
import sys
def next_smallest_palindrome(num):
    numstr = str(num)
    i = num + 1
    while i < sys.maxsize:
        if str(i) == str(i)[::-1]:
            return i
        i += 1
```

This code correctly solves the problem and does exhibit manual counter tracking. However, it is difficult to determine whether a student wrote this code due to holding the misconception or simply because this represents a valid and natural solution approach using a while loop. The LLM-as-a-judge in MCINJECT's validation pipeline tends to be conservative in such cases, labeling these samples as not clearly exhibiting the misconception. While this reduces the exhibit rate, it improves dataset quality by ensuring that retained samples exhibit the misconception more clearly.

This pattern particularly affects operator precedence misconceptions, which achieve approximately 50% exhibit rates. These misconceptions require subtle modifications to arithmetic expressions that can produce correct results or produce unnatural code such as writing `num + 2 * 10` instead of `num + 20`. McInject tends to focus on emulating the code structure rather than writing code that the student with the misconception would write and genuinely believes that it would solve the problem correctly.

A.2.2 Unnatural Code

A second failure mode occurs when MCINJECT produces valid code matching the misconception pattern but lacking semantic plausibility. Successfully injecting subtle misconceptions requires theory of mind, the LLM must authentically adopt the perspective of a student holding that false belief. When this fails, the tool produces code that exhibits the pattern but makes no logical sense.

For instance, for Misconception 18 (*Student believes that in the expression x == a or b, the comparison operator distributes to both operands of the or operator*), MCINJECT generated for the same palindrome problem:

```
import sys
def next_smallest_palindrome(num):
    numstr = str(num)
    for i in range(num+1, sys.maxsize):
        if str(i) == str(i)[::-1] or str(i)[0]:
            return i
```

This code is unrealistic. It makes no sense for a student to check whether `str(i) == str(i)[0]` (checking if a string equals its first character) when they could simply check `len(i) == 1`. Furthermore, the or condition is not needed at all. The program is correct without it. The code syntactically matches the misconception pattern but fails the semantic plausibility test.

Similarly, for (*Student believes that chained function calls are evaluated from right to left*, MCINJECT produced:

```
import sys
def next_smallest_palindrome(num):
    numstr = str(num)
    for i in range(num+1, sys.maxsize):
        if str(i).replace('0', '')[::-1] == str(i).replace('0', ''):
            return i
```

While this attempts to demonstrate order-dependent operations through chained method calls, the code is buggy and nonsensical. Removing zeros from a number before or after reversing is not a correct solution to the palindrome problem and does not reflect a plausible student reasoning path.

B McMiner Analysis

In this section, we examine the performance of McMiner on the benchmark dataset.

B.1 Performance on Correct-Only vs. Misconception Bags

A critical evaluation metric for misconception mining tools is their ability to correctly identify when code contains no misconceptions, avoiding false positives that could mislead educators. Analysis of MCMINER-M performance reveals a gap between these two scenarios across different model configurations, as shown in Table 3.

Model	Reasoning	Correct-Only Accuracy	Misc. Bags Accuracy
Claude Sonnet 4.5	✓	59.30%	89.72%
Claude Sonnet 4.5	×	41.86%	91.30%
OpenAI o3-mini (medium)	–	66.28%	80.63%
OpenAI o3-mini (low)	–	75.58%	75.49%
Gemini 2.5 Flash	✓	36.05%	90.91%
Gemini 2.5 Flash	×	61.63%	60.87%

Table 3: Performance comparison of MCMINER-M across different models on correct-only bags versus bags containing misconceptions. All metrics are based on the 339 total bags (279 bags with misconceptions and 60 correct-only bags).

The results reveal substantial variation in false positive rates across models. Claude Sonnet 4.5 with reasoning achieves 89.72% accuracy on misconception bags but only 59.30% on correct-only bags. Without reasoning, this gap widens (91.30% vs. 41.86%), indicating that reasoning capabilities help reduce false positives. OpenAI o3-mini (low) demonstrates the best balance with nearly equal performance on both bag types (75.58% vs. 75.49%).

B.1.1 Common False Positive Patterns

Analysis of false positives on correct-only bags reveals systematic patterns in the types of "misconceptions" models incorrectly identify **Inefficiencies as Misconceptions**. This is where models sometimes identify "inefficiencies" as misconceptions, such as flagging manual iteration instead of using built-in functions like `enumerate()` or list comprehensions. While these may represent less idiomatic code, they do not constitute misconceptions about language semantics.

B.2 Novel Misconception Discovery

A key strength of MCMINER is its ability to discover misconceptions not present in the predefined misconception bank. Using LLM-as-a-judge validation (Section H), we identified novel true positives—predicted misconceptions that, while not matching ground truth, accurately describe genuine programming misunderstandings exhibited in the code.

Claude Sonnet 4.5 with reasoning discovered 41 novel true positives across the 339 bags (12.1%), while OpenAI o3-mini (medium effort) found 35 (10.3%).

B.2.1 Cherry-Picked Novel Discovery

Example: Division Operator Type Semantics

For a bag containing code implementing tetrahedral number (which is always an integer) calculation, MCMINER identified:

> **Predicted Misconception:** "The student believes that the division operator `/` and integer division operator `//` are interchangeable or that `/` will automatically return an integer when the result is a whole number."
>
> **Code Context:**
>
> ```
> def tetrahedral_number(n):
>  return (n * (n + 1) * (n + 2)) / 6
> ```
>
> **Explanation:** In Python 3, the `/` operator always returns a float, even when dividing two integers that result in a whole number. For calculating tetrahedral numbers, which are always integers, using `//` would be more appropriate to return an integer type.

This misconception was not in the original bank but represents a valid and common misunderstanding about Python 3's division semantics.

B.3 Error Analysis

Certain misconceptions in the benchmark proved particularly challenging for MCMINER, with several achieving near-zero detection rates. Analysis reveals these share a common characteristic: they are mostly *benign misconceptions*, false beliefs that do not cause functional bugs but instead reflect stylistic preferences or suboptimal practices.

Low-Performing Misconceptions:

- "Student believes that the `return` statement requires parentheses around its argument" code like `return(x + y)` is syntactically valid and functions correctly
- "Student believes that function parameters automatically change in recursive calls without explicit modification.". This is a harmful misconception that can cause infinite recursion.
- "Student believes that loop iteration requires manual counter tracking with an explicit variable to access element indices.". This benign misconception can often result in functionally correct code that is also natural and lacks unusual features.

These benign misconceptions fail detection because:

1. **Code is functionally correct**: Models trained on code understanding tasks prioritize functional correctness, making them less sensitive to stylistic variations
2. **Patterns lack distinctive signatures**: Unlike misconceptions that produce syntax errors or logical bugs, benign misconceptions produce code indistinguishable from intentional stylistic choices
3. **Ambiguous intent**: Without access to the student's reasoning, it is impossible to definitively determine whether `return(x)` reflects a misconception or a stylistic preference influenced by other programming languages

C McMiner Interface

To facilitate the use of McMiner for educators and researchers, we developed an interactive web application using the Streamlit[1] Python library that implements the McMiner-S single-instance mining approach. The interface allows users to input a problem description and student code, then select from multiple state-of-the-art LLMs (Claude, GPT, Gemini) to analyze the code for potential programming misconceptions. The tool uses the same prompt template and model configurations as our benchmark experiments, automatically enabling reasoning capabilities for compatible models. Results are displayed with structured misconception descriptions, explanations of how they manifest in the code, and expandable reasoning traces from the model. The application loads user API credentials securely from environment variables and provides real-time analysis of student code within seconds.

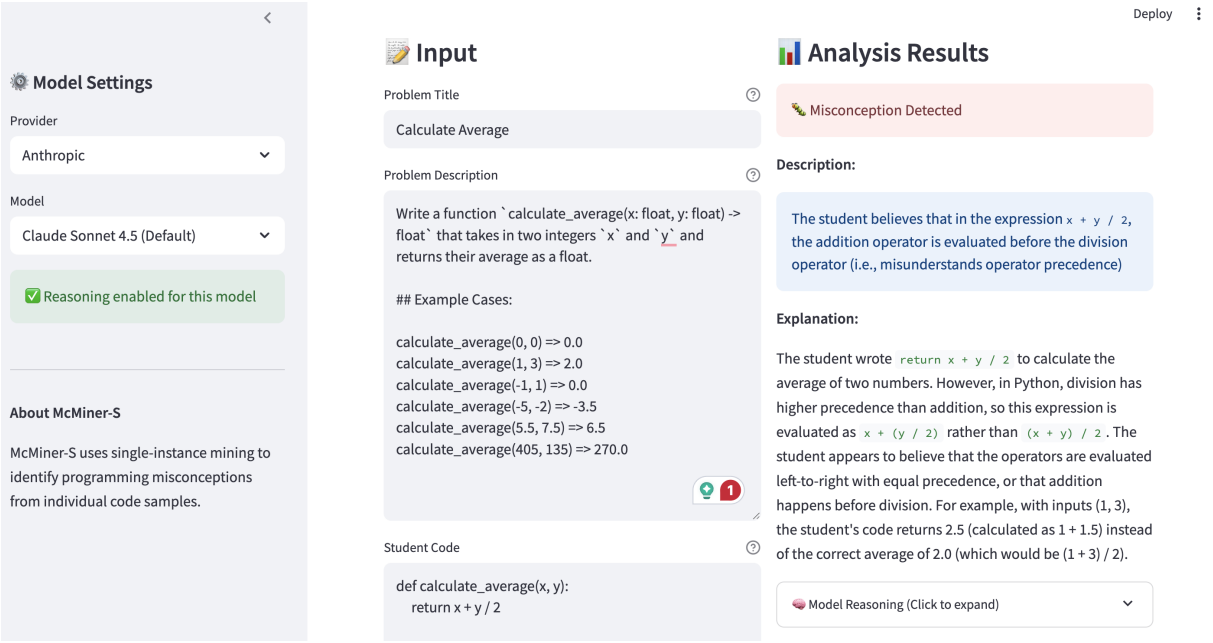

Figure 2: Web application interface of the McMiner-S tool. Users are able to input a programming problem and a student implementation, and analyze the code for any potential misconceptions.

D Dataset Sources

Our misconceptions were adapted from: The Python Misconception bank (Chiodini et al., 2021) (17 misconceptions), SIDELib (Evans et al., 2023) (36 misconceptions), and 14 were composed by the authors, 3 of which are out-of-domain to test the ability of misconception mining tools to identify previously unseen patterns. The three out-of-domain misconceptions are: (1) *Student believes that every variable must be explicitly deleted with 'del' after its last use to prevent memory leaks*, (2) *Student believes that variable names containing vowels (a, e, i, o, u) can only store string values, while consonant-only names can store any type*, and (3) *Student believes that list indexing starts at -1*. While we have 67

[1] https://streamlit.io

misconceptions, our benchmark contains 1,063 code samples exhibiting those misconceptions. Our benchmark can be easily extended given any new misconception.

The programming problems, solutions, and unit tests in our benchmark dataset were curated from multiple established sources to ensure diversity and representativeness. Each problem includes comprehensive test cases, with an average of 3.31 unit tests per problem to validate correctness. The majority of problems (438 out of 501, or 87.4%) were adapted from the MBPP dataset (Austin et al., 2021), which provides a comprehensive collection of simple computational problems in Python with corresponding solutions and test cases. To supplement this core set, we incorporated 27 problems from the Socratic Debugging dataset (Al-Hossami et al., 2023, 2024), which focuses on debugging scenarios, along with 8 problems from the Auckland dataset (Ettles et al., 2018) and 6 from the FalconCode dataset (de Freitas et al., 2023). Finally, 19 problems were handwritten by the authors, drawing inspiration from public programming courses and educational websites such as w3resource to fill gaps in programming language construct coverage.

Table 1 summarizes the key statistics of our benchmark dataset. The dataset was generated from 67 distinct misconceptions applied to 25 randomly sampled problem-solution pairs ($N_k = 25$) from our curated set of 501 problems. The MCINJECT tool generated 1,675 code samples, of which 498 were deemed inapplicable due to incompatibility between the misconception and the problem-solution pair, resulting in 1,177 generated code samples. These samples were then validated using LLM-as-a-judge evaluation, which filtered out 114 samples that did not properly exhibit their intended misconceptions. This yields 1,063 code samples that successfully exhibit misconceptions. Code samples that were deemed inapplicable or did not properly exhibit their intended misconceptions were replaced with the original correct solution and have a null misconception. This yields a final dataset of 1,675 code samples. To enable evaluation in a Multiple Instance Learning setting (later described in Section 4), we organized these samples into 339 bags, each containing 4-8 code samples. This bag structure allows us to evaluate both single-instance mining (MCMINER-S) and multi-instance mining (MCMINER-M). Importantly, 60 bags (17% of the bags) contain only correct code samples without any misconceptions, serving as negative examples to test the system's ability to correctly identify when no misconceptions are present.

D.1 Misconception Description Writing Guidelines

All misconceptions (composed, adapted, and ooD) were written by hand and include one example of a code snippet exhibiting that misconception. Figure 3 shows the writing guidelines used to create the misconception bank.

Operational definition: A concise, one-sentence statement of the student's incorrect belief about a programming construct, operation, or API, excluding the correct concept.

Guidelines

- **State ONLY the incorrect belief**: Use "Student believes that..." without mentioning or implying the correct construct, concept, or fix.
- **Avoid line numbers**: Reference only the operation or built-in/API involved, not specific line numbers.
- **Be concrete and specific about the construct**: Mention the relevant operation, method, or construct (e.g., `list.pop`, string indexing, `range` bounds, conditional statements).
- **Focus on ONLY one construct**: Write the misconception in terms of only one construct, operation, or API. Avoid general misconceptions applicable to multiple constructs.
- **Generalize the misconception**: The misconception should be applicable to multiple implementations, not specific to a single problem or implementation.
- **Exclude inputs, outputs, and fixes**: Do not include example inputs, expected/actual outputs, or the bug fix.
- **Keep it concise**: One clear sentence that maps directly to what the student implemented.

Figure 3: Misconception description writing guidelines.

E The McInject Tool

The MCINJECT tool leverages Claude-Sonnet-4.5 with extended thinking enabled to implement the misconception injection task. We use temperature 1.0, max_tokens 6000 (4000 output + 2000 thinking budget), and a 2000 token thinking budget. The tool uses zero-shot prompting with structured XML output format, where the LLM is instructed to modify the correct solution such that it appears as if written by a student who genuinely believes their approach is correct, despite holding the specified misconception.

To enhance generation quality, the prompt includes illustrative examples alongside misconception descriptions. Each misconception in the dataset includes a concrete example demonstrating how the false belief manifests in code, which helps guide MCINJECT to produce more realistic student-like code. The McInject prompt template is shown in Figure 4.

To further improve code quality, we implement an iterative refinement process. After MCINJECT generates code, an LLM-as-a-judge evaluates whether the code properly exhibits the target misconception (see Section H.1). If the judge provides actionable feedback indicating the misconception is not clearly exhibited, this feedback is appended to the original conversation context, and MCINJECT is prompted to refine the code. This feedback loop can iterates once.

Your Task
You will be given as input a programming *problem*, *code* implementing a solution, and the description of a *misconception*. A misconception is a false belief that the student holds about some programming language construct. Note that, misconceptions do not always result in buggy code. Some misconceptions lead to stylistic differences or inefficiencies rather than errors. For example, a student who believes "all variable identifiers must use only one letter" might write working but less readable code.
Modify the code such that it looks as if written by a student who has that misconception. The student genuinely believes that this modified code solves the problem correctly.

Input Format
Problem Description: [problem_description]
Given Implementation:

```
[correct_solution]
```

Misconception Description:

```
[misconception_description]

**Example:**
[misconception_example]
```

Output Format

```
<code>
[The complete modified Python code exhibiting the
misconception]
</code>
```

If the misconception relates to language constructs that are not present in the given solution, output:

```
<code>
NONE
</code>
```

Figure 4: Prompt template for MCINJECT tool.

A key feature of MCINJECT is its ability to handle incompatible cases where a misconception cannot be meaningfully applied to a given solution. Misconceptions are considered incompatible when they relate to language constructs or concepts that are not present or relevant in the solution. For example, a misconception about loop behavior would be incompatible with a solution that contains no loops. In such cases, the tool indicates inapplicability rather than forcing inappropriate modifications.

All code samples produced by MCINJECT undergo post-processing where inline comments are automatically removed to ensure clean output. This process uses Python's tokenize module and includes syntax validation to maintain code correctness.

To validate the quality of generated code samples by McInject, we employed LLM-as-a-Judge evaluation showing that 90.3% of the code samples (1,063 out of 1,177 generated samples) successfully

exhibit their target misconceptions. We also employ manual evaluation of 38 samples composed of a misconception and a problem-solution pair that were deemed to be inapplicable by McInject, and observe 100% agreement with human judgment.

We note that this is a conservative estimate of McInject's performance, since the LLM-as-a-judge discards code samples that do exhibit the misconception, yet are both correct and natural, making the misconception extremely difficult to detect. Manual evaluation of 88 code samples demonstrates 96.6% agreement between LLM judgment and human assessment, confirming the reliability of the LLM-as-a-judge approach in determining whether a code sample exhibits a misconception description. The LLM-as-a-judge leveraged Claude-Sonnet-4.5 with the same hyperparameters as McInject.

F The McMiner Tools

This section lists all the prompts that are used by the McMiner tools.

Key Terminology

A **programming misconception** refers to a false belief that a student holds about some programming language construct or built-in function in Python. Programming misconceptions can be about the syntax or the semantics of constructs in the Python programming language. They should not be about concepts in the problem description.

Your Task
Given a problem description and student code that attempts to solve that problem, identify a most likely programming misconception that is exhibited by that code, if any.

Input Format
Problem Description: {problem_description}
Student Code:

```
{student_code}
```

Output Format

```
<misconception>
<description>[Describe misconception, starting
with "The student believes"]</description>
<explanation>[Explain how the code exhibits
          the misconception]</explanation>
</misconception>
```

If no misconceptions are found, output:

```
<misconception>NONE</misconception>
```

Figure 5: Prompt template for MCMINER-S (single-instance mining). The full template includes additional metadata fields and guidelines.

G Experimental Setup

We evaluated MCMINER using three state-of-the-art LLMs through their respective APIs:

- **OpenAI**: We used the o3-mini model with two reasoning effort levels. For low effort, we set reasoning_effort to "low" and max_completion_tokens to 7000 (4000 output + 3000 reasoning). For medium effort, we set reasoning_effort to "medium" and max_completion_tokens to 9000 (4000 output + 5000 reasoning). Note that o-series models do not support temperature configuration.

- **Anthropic**: We used the claude-sonnet-4-5 model. For non-reasoning experiments, we used temperature 0.1 and max_tokens 4000. For reasoning-enabled experiments, we used temperature 1.0 (as required by Anthropic for extended thinking), max_tokens 6000 (4000 output + 2000 thinking budget), and enabled extended thinking with a 2000 token budget.

- **Google Gemini**: We used the gemini-2.5-flash model with temperature 0.1 and max_tokens 4000. For reasoning-enabled experiments, we increased max_tokens to 6000 (4000 output + 2000 thinking budget) and enabled thinking with a 2000 token budget.

Key Terminology

A **programming misconception** refers to a false belief that a student holds about some programming language construct or built-in function in Python. Programming misconceptions can be about the syntax or the semantics of constructs in the Python programming language. They should not be about concepts in the problem description. For example, "Student believes range(n) produces values from 1 to n inclusive" is a valid programming misconception about Python's range() function, while "Student thinks natural numbers can be negative" is not a programming misconception. Also, while misconceptions often cause bugs, sometimes they do not. For example, "Student believes all variable names need to have exactly one letter" is a programming misconception that does not necessarily cause a bug.

Your Task
Input: a set of <problem description, code> pairs, where the code in each pair attempts to solve the corresponding problem.
Output: the description of a programming misconception that is exhibited by one or more code samples in the input set. While the input set will often contain code samples that show no misconception, try to identify a misconception that is exhibited by most code samples in the input set. If you cannot find any misconception, output NONE.

Input Format
Problem Description: {problem_description}
Student Code:

```
{student_code}
```

Output Format

```
<misconception>
<description>[Describe misconception, starting
with "The student believes"]</description>
<explanation>[Explain how the code exhibits
          the misconception]</explanation>
</misconception>
```

If no misconceptions are found, output:

```
<misconception>NONE</misconception>
```

Figure 6: Prompt template for MCMINER-M (multi-instance mining). The full template includes additional metadata fields and guidelines.

All experiments used individual request processing (no batching) with a zero-shot prompting strategy. For evaluation, we used the benchmark dataset $\mathcal{D}$ described in Section 4, containing the 339 bags of code samples.

To evaluate whether predicted misconceptions semantically match ground truth (GT) misconceptions, we employed an LLM-as-judge approach using Claude Sonnet-4.5 with reasoning capabilities. The evaluation uses identical hyperparameters as MCINJECT. Claude evaluates whether the predicted misconception description matches the intended misconception.

To credit novel misconception discoveries not present in our benchmark, we implement a novelty-aware evaluation approach. For predictions that do not match GT, we use the same LLM-as-judge validation from MCINJECT to determine whether the code samples accurately exhibit the predicted misconception. If validated, these are counted as novel true positives.

H Evaluation Prompts

This section discusses all the prompts used for evaluation in this paper.

H.1 McInject LLM-as-a-judge Prompt

To validate that code samples generated by MCINJECT properly exhibit their intended misconceptions, we employ an LLM-as-a-judge evaluation. The judge is given a misconception description with an example, and a code sample to analyze. The task is to determine whether the code exhibits the misconception, providing a binary judgment (Y/N) and optional feedback.

When used in the iterative refinement process, the feedback mechanism works as follows: if the judge determines the misconception is not clearly exhibited and provides specific feedback (i.e., feedback is not "NONE"), this feedback is used to create a multi-turn conversation. The original MCINJECT prompt and its initial response are preserved, and the feedback is appended as a new user message instructing

MCINJECT to improve the code based on the critique. This creates a conversational refinement loop where MCINJECT can iteratively improve its output. The process terminates when either the judge confirms the misconception is exhibited (feedback is "NONE"), or the iteration limit is reached.

Importantly, the prompt emphasizes that misconceptions do not necessarily cause bugs—code can be syntactically and logically correct while still exhibiting a false belief pattern. Figure 7 shows the prompt template used for this evaluation.

Input Format
The Misconception

```
<misconception>
Description: {misconception_description}
Example: {misconception_example}
</misconception>
```

The Code to Analyze

```
<code>
{code_to_analyze}
</code>
```

Your Task
Determine whether this code exhibits the misconception described above.

Key Understanding
A misconception does NOT necessarily induce bugs or errors! Code can be:

- **Syntactically correct** (no syntax errors)
- **Logically correct** (produces expected output)
- **Yet still exhibit a misconception** (shows the student holds a false belief)

Analysis Guidelines

1. **Understand the misconception deeply**: What incorrect belief does the student have? What coding patterns would reveal this belief?
2. **Analyze the code systematically**: Look for patterns that match the misconception. Check if the code structure reflects the incorrect belief.
3. **Focus on the belief, not the outcome**: Does the code structure suggest the student holds this false belief? Even if the code works, does it show the misconception pattern?

Output Format

```
<answer>
<exhibits_misconception>Y or N</exhibits_misconception>
<feedback>[Optional feedback]</feedback>
</answer>
```

Figure 7: Prompt template for LLM-as-a-judge evaluation of MCINJECT generated code samples. The full prompt template includes metadata fields to output such as rationale and confidence level.

H.2 McMiner Semantic Matching Prompt

To evaluate MCMINER predictions, we use an LLM-as-a-judge to determine whether a predicted misconception semantically matches the ground truth misconception. The judge is provided with the ground truth misconception, the predicted misconception, and the code samples that were analyzed. The task is to assess whether both descriptions capture the same fundamental programming misunderstanding, accounting for natural variations in wording. The evaluation considers core concept alignment, evidence in the code samples, and semantic equivalence. Figure 8 shows the prompt template used for this evaluation.

Input Format
Ground Truth Misconception

```
{ground_truth}
```

Predicted Misconception

```
{predicted_misconception}
```

Code Samples Analyzed

```
{code_samples}
```

Task
Determine if the predicted misconception accurately captures the same conceptual misunderstanding as the ground truth. Consider:

1. **Core Concept Match**: Are they describing the same fundamental misunderstanding about programming concepts?
2. **Evidence Alignment**: Does the predicted description align with what's shown in the code samples?
3. **Semantic Equivalence**: Minor wording differences are acceptable if the core concept matches.
4. **Example Reference**: If provided, use the misconception example to better understand the typical manifestation of this misconception.

Special Cases

- If ground truth is "NO MISCONCEPTION" and prediction found no misconceptions, this is a match.
- If ground truth is "NO MISCONCEPTION" but prediction found misconceptions, this is NOT a match.
- If ground truth describes a misconception but prediction found none, this is NOT a match.

Output Format

```
<evaluation>
<match>true or false</match>
</evaluation>
```

Figure 8: Prompt template for semantic matching evaluation of McMiner predictions. The full prompt template includes metadata fields to output such as confidence level and explanation.