

Multimodal Programming in Computer Science with Interactive Assistance Powered by a Large Language Model

Rajan Das Gupta¹, Md. Tanzib Hosain¹, M. F. Mridha¹, and Salah Uddin Ahmed²

¹American International University-Bangladesh

²University of South-Eastern Norway



Introduction

Problem: Large introductory computer science courses struggle to provide students with timely and effective homework assistance.

Opportunity: Large Language Models (LLMs) like DeepSeek R1 can offer instant, interactive support.

Challenge: How can we use LLMs to facilitate learning without simply giving students the answers?

Our Solution: An interactive homework help system integrated into the students' programming environment that provides pedagogical guidance rather than direct solutions.

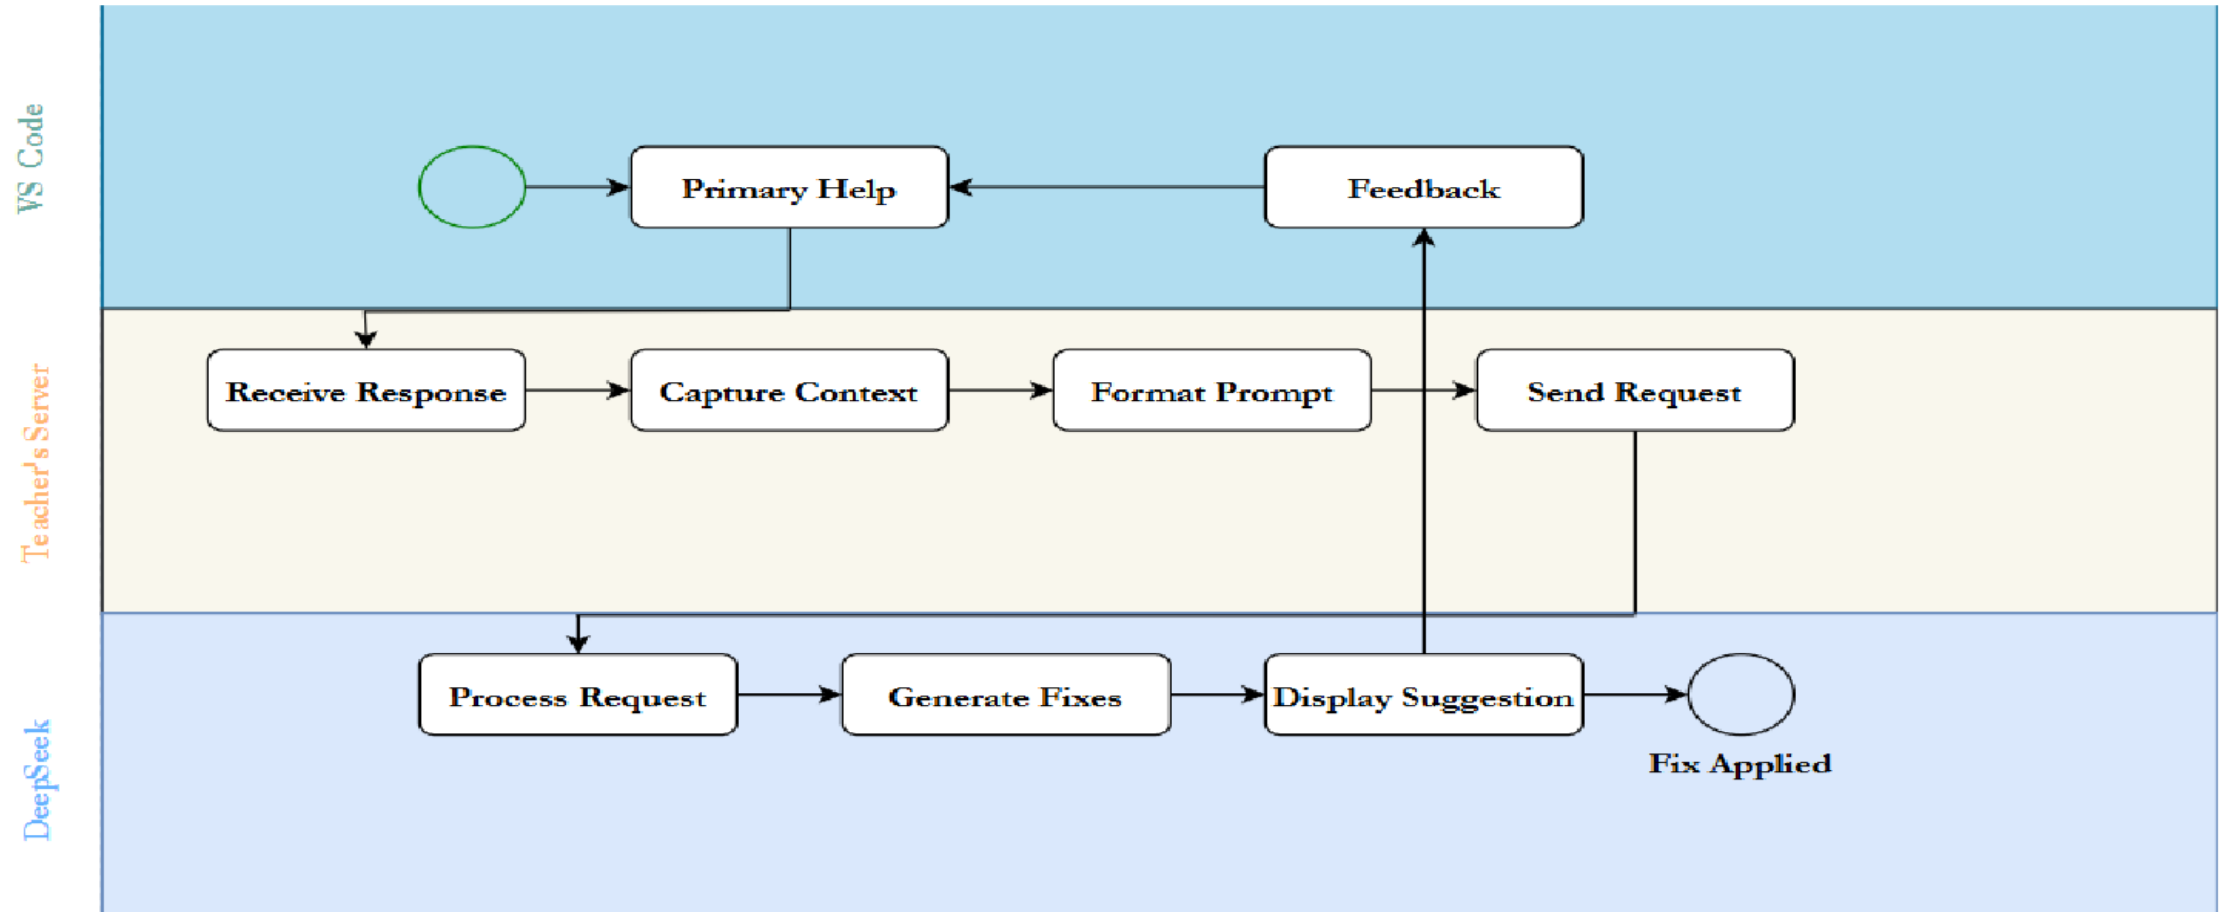
System Overview

Our system provides help through two main channels, both powered by the same backend assistant.

VS Code Extension: A "Primary Help" button directly within the code editor.

Command-Line Auto evaluator: A "Feedback" option integrated into the tool students use to check their code.

System Overview



Methodology

We decided against fine-tuning a model and instead focused on **prompt engineering** to guide the DeepSeek R1 model.

Goal: To provide feedback that mirrors how a human tutor would help.

Our Prompt's Strategy:

- Assess the student's understanding of the problem.
- Identify conceptual gaps.
- Help the student formulate a plan.
- Provide hints or code templates as a last resort, without giving away the solution.

The "Secret Sauce": Our Prompt Design

We developed a detailed prompt that instructs the LLM on how to behave like an effective tutor.

Key Instructions to the LLM:

"Your role is to guide students in learning programming concepts effectively."

"Avoid providing direct answers or complete code."

"If the code is incomplete or incorrect, consider the following steps..." (e.g., "Does the student lack any fundamental understanding?", "Is the existing code moving in the right direction?")

"Use a Socratic approach to encourage critical thinking."

"Keep your response brief - one or two sentences at most."

Challenges & Solutions

Challenge: The LLM often gave direct answers, especially in longer conversations.

Solution: We restricted the interaction to a "one-shot" request to limit this behavior.

Challenge: The LLM would sometimes "correct" perfectly good code (a false negative).

Solution: We experimented with having the LLM generate a correct solution first for comparison, but this added too much latency. The issue remains a key area for improvement.

Challenge: The LLM would suggest solutions that were technically correct but violated the specific constraints of the assignment.

Solution: We added assignment-specific context to the prompt.

Deployment & Findings

Rollout: Deployed in two phases to over 1100 students in an introductory programming course.

Positive Findings:


- Students used the bot frequently, especially near deadlines, suggesting they found it helpful.
- The bot was effective at helping students resolve syntax errors and function misuse.
- A noticeable decrease in the number of support inquiries during office hours was observed.

Negative Findings:

False Negatives: The bot would sometimes claim correct code was wrong.

False Positives: The bot would approve incorrect code or provide misleading suggestions, leading to student frustration.

Deployment & Findings



Limitations: The "Failure" Patterns

The most significant issue was the bot's unreliability, which manifested in two main ways:

False Negatives:

- The bot tells a student their correct code is incorrect.
- This is less harmful in the auto evaluator mode, as passing all test cases overrides the bot's feedback.

False Positives (More Damaging):

- The bot approves incorrect code, leading students down a rabbit hole of debugging.
- The bot suggests a fix that violates assignment constraints, causing conflicts with the autoevaluator. This was the most frustrating experience for students.

Conclusions & Future Directions

Conclusion: An LLM-powered assistant can be a valuable tool in programming education, but its effectiveness is highly dependent on careful prompt engineering and managing its limitations.

Future Work:

Improve Accuracy: Continue to refine prompts and potentially explore fine-tuning to reduce errors.

Enhance Interaction: Allow students to ask for clarification (e.g., "Explain that differently") and rate the helpfulness of the advice.

Deeper Analysis:

- How does bot usage correlate with student performance?
- When and why do students seek help from the bot?
- How do students feel about interacting with a bot versus a human TA?

Thank You

Any Questions?