

Poster Session at Graduate School Information Fair

Enhancing Programming Learning with LLMs: Prompt Engineering and Flipped Interaction

Background

Motivation


Education is becoming more and more virtual, with technologies such as eLearning and remote lectures increasing in popularity. However, virtual study poses the challenge of obtaining good, constructive feedback. Instructors can be busy or slow to answer, and some students may be shy to ask questions.

What about ChatGPT?


Recent advances in artificial intelligence and Large Language Models (LLMs) can give students a way to ask for feedback and get immediate results. However, this method may be too easy for students: while an instructor may help a student to learn, an LLM is likely to simply give student raw answers.

Goal

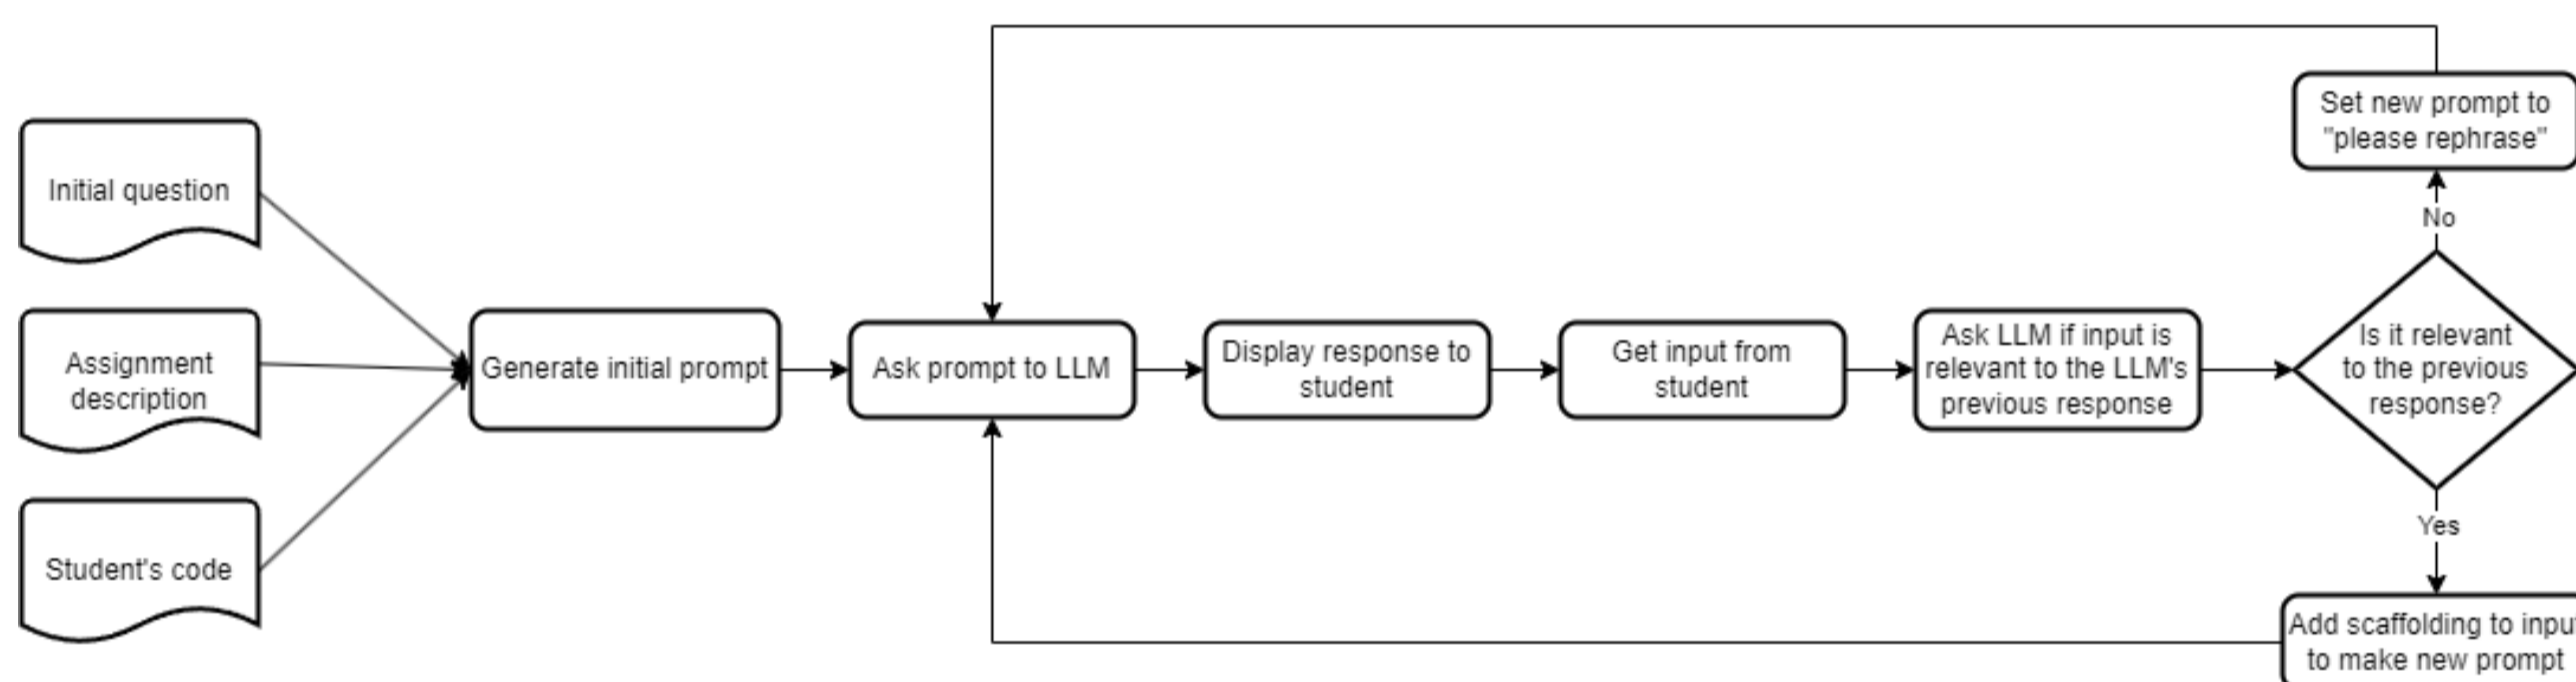
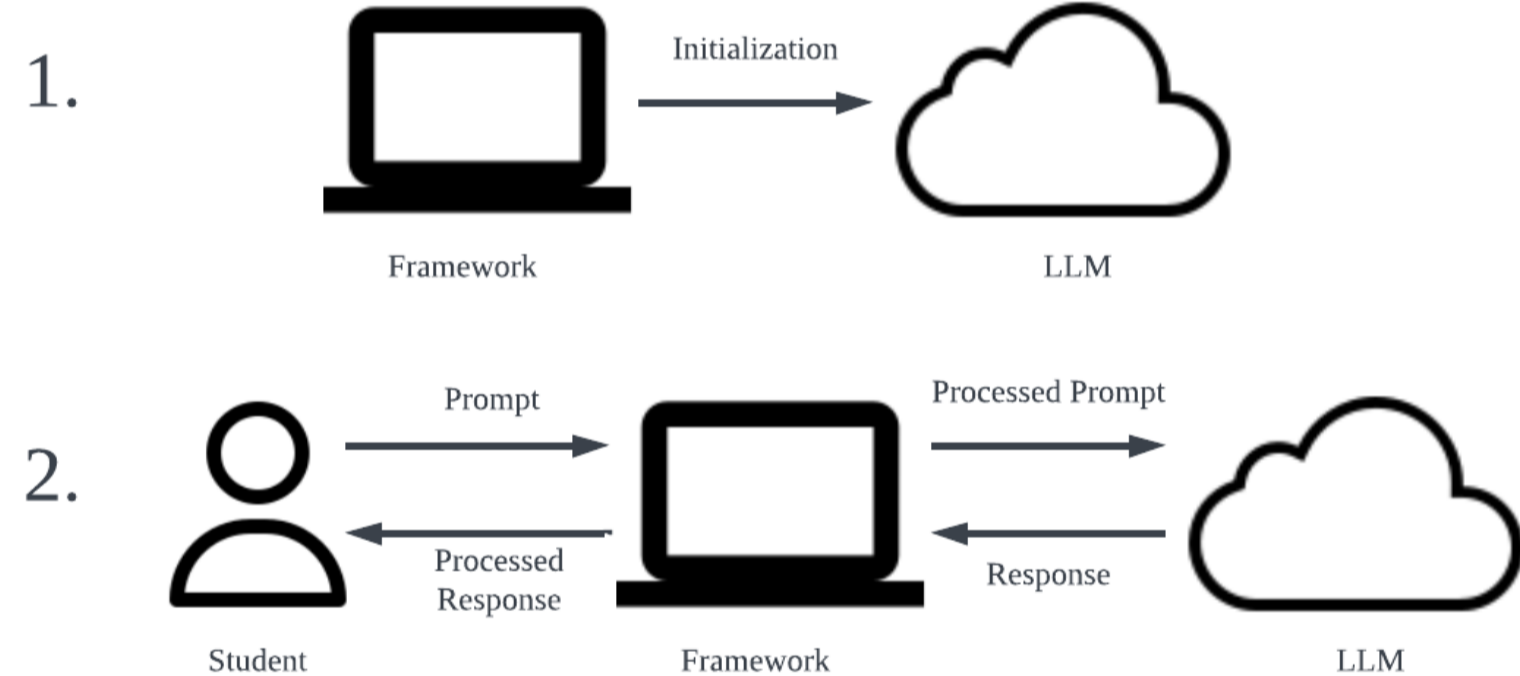
The goal of this research is to develop and test a way to use LLMs to assist students with their programming assignments in a way that is conducive to their learning.



<https://www.polarismarketresearch.com/industry-analysis/e-learning-market>



Methodology

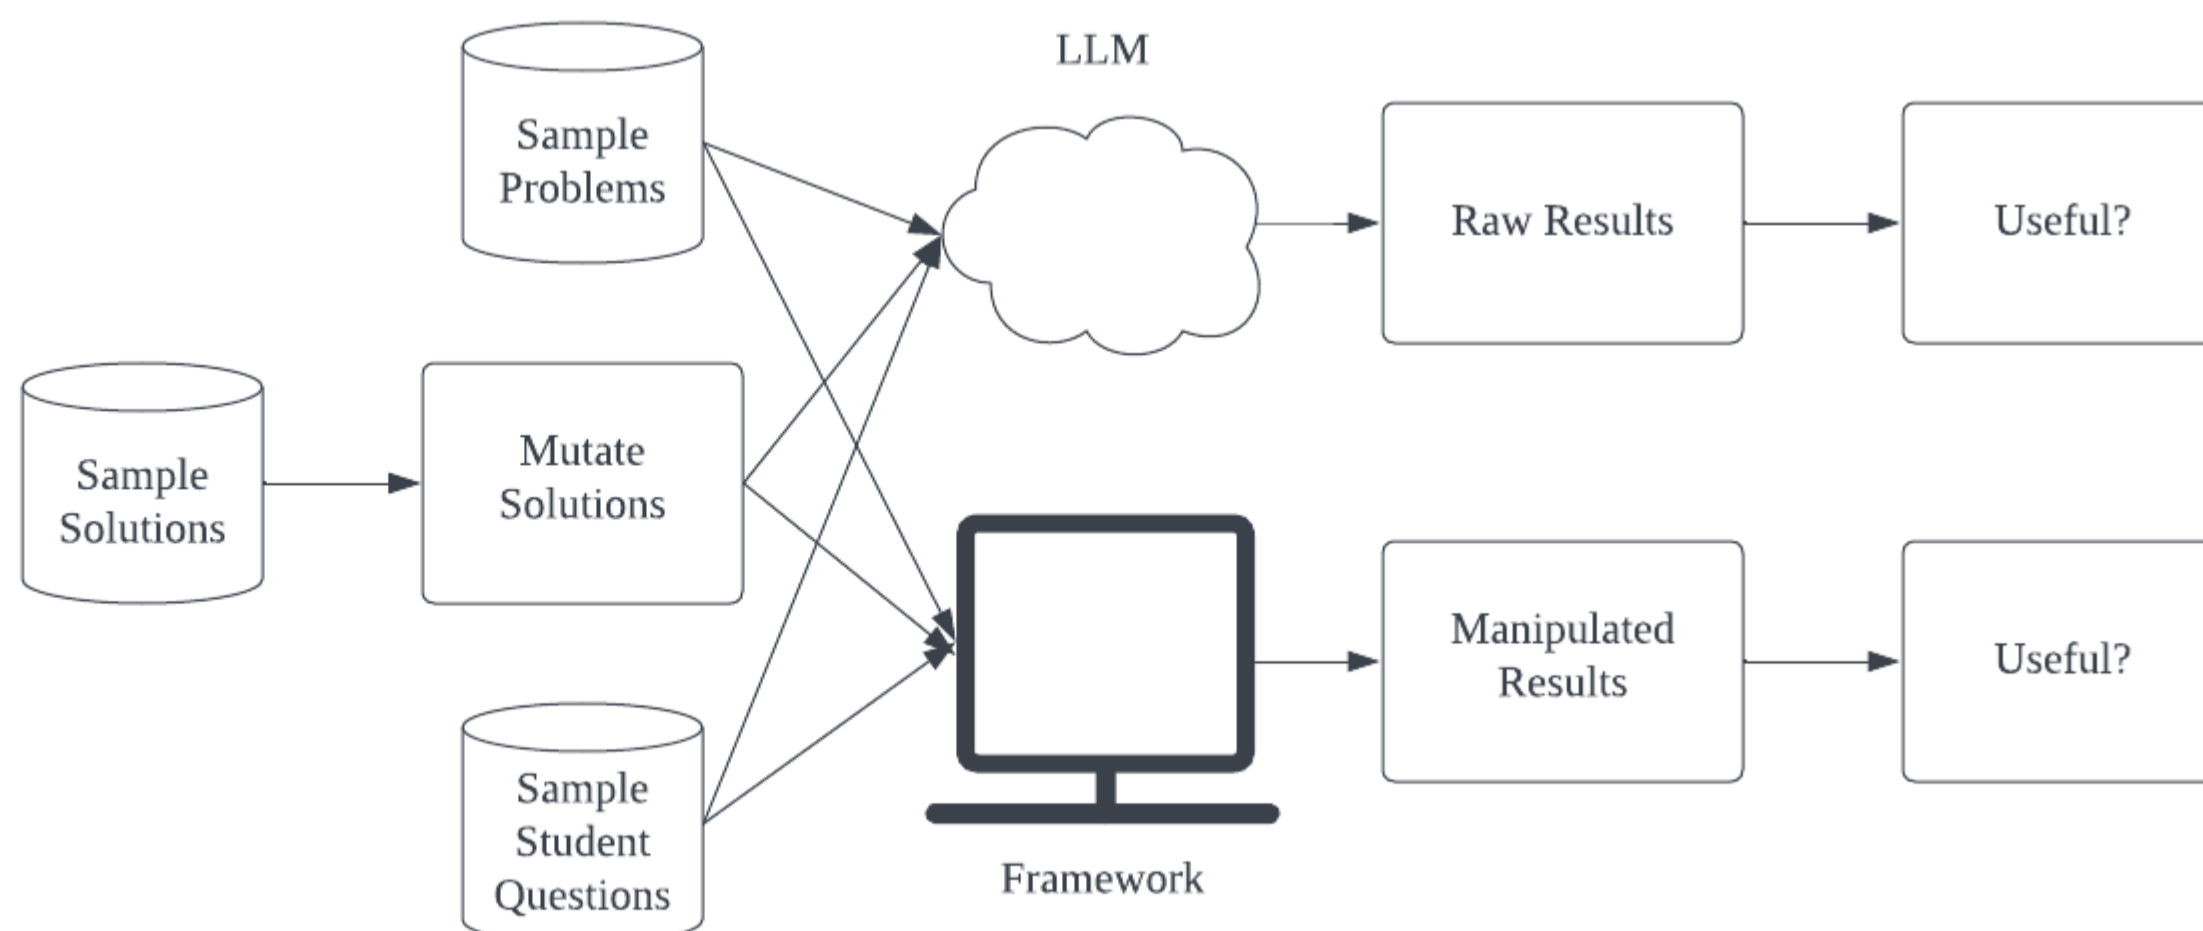


We designed a framework intended to be implemented as an IDE plugin. The student initializes a conversation with the framework from within the IDE, then their code, an assignment description, and an initial prompt from the student are sent to the framework. The framework communicates with an LLM and outputs a refined response to the student.

Results

Experiment Setup

We took sample problems and solutions from geeksforgeeks.com, mutated the solutions so that they mimicked what a student might write, and fed those into both the framework and raw LLM. We then used ChatGPT to generate responses for the student's side of a conversation with the LLM or Framework, and compared the framework's responses to the raw LLM's responses based on two metrics.




Overall Conversation

We measured the perceived usefulness of overall conversations. A conversation was considered a pass if it appears to be pushing the student in the right direction without giving any exact answers. In the table below, the Framework Case refers to conversations with our framework, and the Raw Case refers to conversations with LLM alone.

	Passes	Fails
Framework Case	10	4
Raw Case	0	14

Individual responses

We measured the perceived usefulness of individual responses. Here we categorize each response from the framework as Failure (appeared in a failed conversation from the other experiment), Useful (could push the student in the right direction), Repeated Response, Exact Answer, and Misleading.



Poster Session at Graduate School Information Fair Enhancing Programming Learning with LLMs: Prompt Engineering and Flipped Interaction

Background

Motivation


Education is becoming more and more virtual, with technologies such as eLearning and remote lectures increasing in popularity. However, virtual study poses the challenge of obtaining good, constructive feedback. Instructors can be busy or slow to answer, and some students may be shy to ask questions.

What about ChatGPT?

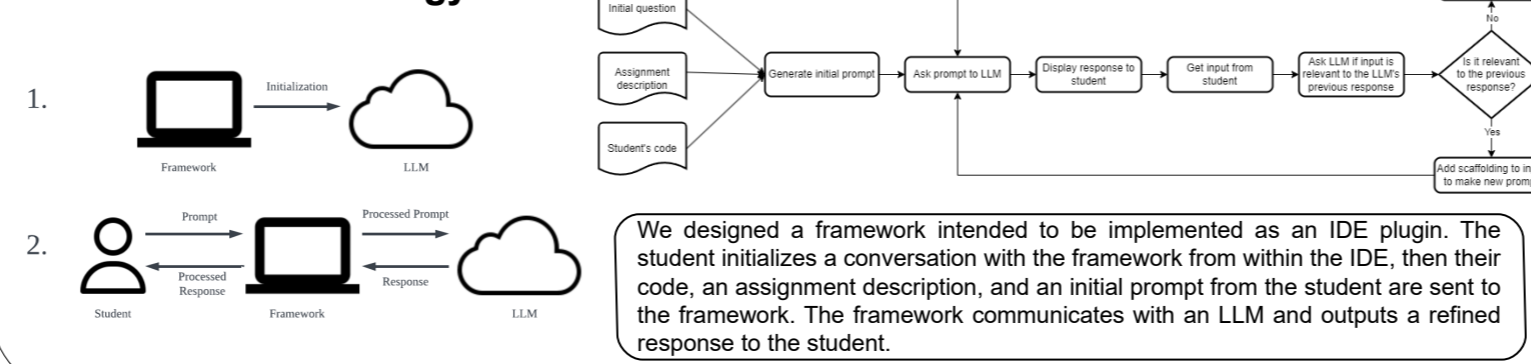
Recent advances in artificial intelligence and Large Language Models (LLMs) can give students a way to ask for feedback and get immediate results. However, this method may be too easy for students: while an instructor may help a student to learn, an LLM is likely to simply give student raw answers.

Goal

The goal of this research is to develop and test a way to use LLMs to assist students with their programming assignments in a way that is conducive to their learning.



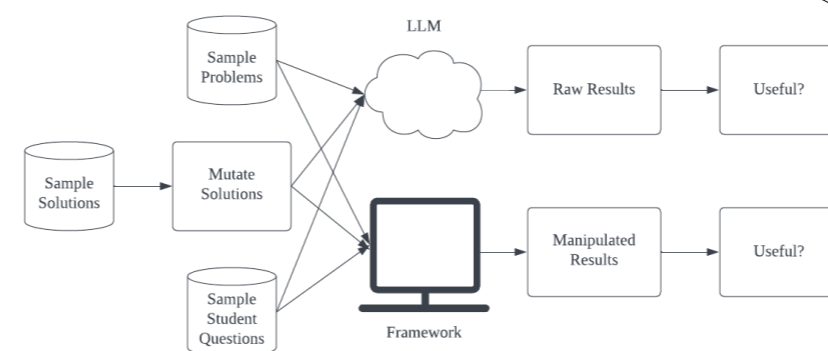
Methodology



Results

Experiment Setup

We took sample problems and solutions from geeksforgeeks.com, mutated the solutions so that they mimicked what a student might write, and fed those into both the framework and raw LLM. We then used ChatGPT to generate responses for the student's side of a conversation with the LLM or Framework, and compared the framework's responses to the raw LLM's responses based on two metrics.



Overall Conversation

We measured the perceived usefulness of overall conversations. A conversation was considered a pass if it appears to be pushing the student in the right direction without giving any exact answers. In the table below, the Framework Case refers to conversations with our framework, and the Raw Case refers to conversations with LLM alone.

	Passes	Fails
Framework Case	10	4
Raw Case	0	14

Individual responses

We measured the perceived usefulness of individual responses. Here we categorize each response from the framework as Failure (appeared in a failed conversation from the other experiment), Useful (could push the student in the right direction), Repeated Response, Exact Answer, and Misleading.

