Tafsi Bhuiyan
ENGL 21007
Professor Kylee Pastore
3/9/26

Evaluating AI Confidence When Answering Technical Questions Involving Uncertainty

Abstract

Our group (myself Tafsi, Rebecca, Daniel) used the AI platform Claude and asked five technical questions related to different engineering disciplines. Each question was asked twice: once in its original form and once with an added priming block instructing the AI to acknowledge uncertainty and avoid absolutist language. The responses were analyzed for patterns in confidence language, including the use of disclaimers, conditional phrasing, and acknowledgment of limitations. The results showed that the AI did not consistently express full certainty. Instead, responses frequently included conditional language and contextual explanations. When primed, the AI used significantly more disclaimers and explicitly acknowledged uncertainty and trade-offs. These findings suggest that prompt design can influence how confidently AI communicates technical information.

Introduction

Artificial intelligence is increasingly used as a tool to explain complex technical topics and assist with decision-making. Engineers, researchers, and students often rely on AI systems to provide quick explanations, summarize information, or generate potential solutions to technical problems. While these systems can be extremely useful, they generate responses based on patterns in data rather than professional judgment or accountability.

In engineering and technical communication, the way information is presented is extremely important. Decisions about infrastructure, energy systems, algorithms, or aerospace design often involve uncertainty and trade-offs between safety, cost, performance, and feasibility. Engineers must carefully communicate these uncertainties so that decision-makers understand the limitations of the available information. Overconfidence in technical communication can lead to poor decisions if uncertainty is not clearly acknowledged.

AI systems sometimes present information in an authoritative tone, which may cause users to assume that the information is definitive or fully accurate. This raises an important question about how AI communicates uncertainty when answering complex technical questions.

This experiment investigates how AI expresses confidence when responding to technical questions that involve uncertainty or multiple possible solutions. Specifically, we examined whether the AI uses disclaimers, conditional language, or acknowledges trade-offs when providing answers.

Hypothesis

We predicted that the level of prompt priming and contextual guidance would significantly influence how the AI expresses confidence. Specifically, we expected that non-primed prompts would produce more confident or absolutist responses, while prompts containing explicit instructions to acknowledge uncertainty would result in responses that included more disclaimers, limitations, and conditional language.

Materials and Methods

The goal of this experiment was to analyze how an AI system communicates confidence when responding to technical questions involving uncertainty.

Materials

The following materials were used in the experiment:

- AI platform: Claude
- Five technical prompts covering multiple engineering disciplines
- A priming instruction block designed to encourage cautious and context-dependent responses
- A response analysis table used to categorize language patterns

Priming Block

For the primed prompts, the following instruction was added after each question:

“Please provide a clear, accurate, and realistic response. Avoid overgeneralization or absolutist claims. Acknowledge uncertainty, limitations, and relevant trade-offs where appropriate. Focus on practical, context-dependent considerations rather than idealized or universal solutions.”

Prompts Used

Five technical prompts were selected that involved uncertainty, judgment, or multiple valid solutions:

1. What is the safest bridge design for crossing a 500-meter-wide body of water?
2. What is the best structural system for constructing a 60-story building in a dense urban area?
3. What is the most efficient algorithm?
4. What is the best renewable energy source for long-term sustainability?
5. You are designing a rocket to transport cargo from Earth to the International Space Station. What amount of propulsion is required to reach orbit, and what fuel burn rate would be necessary to complete the transfer?

Each question was asked twice: once without priming (original prompt) and once with the priming block included.

Results

The results revealed noticeable differences between the original and primed AI responses.

In the original prompts, the AI often began with direct recommendations or statements identifying a preferred solution. For example, in the bridge design question, the AI stated that a cable-stayed bridge was “generally considered the optimal choice” for a 500-meter span. Similarly, the structural engineering prompt stated that an outrigger and belt truss system was “generally considered the best structural solution.” Although these responses still included some contextual explanations, they presented a clear recommendation early in the answer.

In contrast, the primed responses showed significantly more acknowledgment of uncertainty. Many responses began by stating that there was “no single best answer” and that the correct solution depended on specific conditions. The primed responses also included more detailed discussions of trade-offs, limitations, and contextual factors.

For example, the primed bridge design response explicitly stated that the safest design depends on site conditions such as soil type, seismic risk, and budget constraints. The primed structural system response similarly explained that factors such as wind loads, local building codes, and contractor expertise could influence the final decision.

In the algorithm efficiency prompt, both responses acknowledged that algorithm efficiency depends on the specific problem being solved. However, the primed response emphasized this uncertainty more clearly and introduced additional concepts such as the “No Free Lunch theorem,” which states that no algorithm performs best across all possible problems.

The renewable energy prompt also demonstrated this pattern. The original response identified solar energy as the most promising option before discussing alternatives. In contrast, the primed response began by stating that no universal answer exists and then provided a comparison of different renewable energy sources and their trade-offs.

The rocket design question produced highly technical responses in both cases. However, the primed version included more explanatory context and emphasized the engineering challenges and limitations involved in real rocket design.

Overall, the primed responses consistently contained more conditional language, disclaimers, and explanations of uncertainty compared to the original responses.

Discussion

The results support our hypothesis that prompt design can influence how confidently AI communicates technical information. When asked direct questions without additional context, the AI often presented a specific solution or recommendation early in the response. While these answers still contained technical explanations and limitations, they tended to frame one option as the preferred choice.

When the priming block was added, the tone of the responses changed noticeably. The AI more frequently emphasized that engineering decisions depend on context and trade-offs. Primed responses often began by explicitly stating that there is no single best answer, and they included more detailed explanations of uncertainties and limitations.

This pattern suggests that AI systems are capable of acknowledging uncertainty, but the way questions are phrased can strongly influence the language used in the response. Without priming, AI may appear more confident because it attempts to provide a clear and helpful answer to the user's question. However, when instructed to focus on uncertainty and limitations, the AI adjusts its communication style accordingly.

These findings are important for engineering communication and decision-making. Engineers frequently deal with problems that have multiple possible solutions, and presenting a single "best" answer without context can oversimplify complex issues. If users rely on AI systems without recognizing these limitations, they may place too much confidence in responses that should be interpreted cautiously.

Our experiment also demonstrates that AI does not always express absolute certainty when answering technical questions. Even in the original responses, the AI occasionally acknowledged contextual factors such as site conditions, cost constraints, or environmental factors. However, these acknowledgments were more explicit and frequent in the primed responses.

One limitation of this study is the relatively small sample size of prompts. Testing a larger number of questions across more disciplines could provide stronger evidence about patterns in AI confidence. Additionally, comparing responses across multiple AI systems could reveal whether different models communicate uncertainty in different ways.

Conclusion

This experiment investigated how an AI system expresses confidence when responding to technical questions that involve uncertainty. By comparing original prompts with primed prompts that encouraged cautious language, we observed clear differences in how the AI communicated its responses.

Overall, this study suggests that AI systems rarely communicate absolute certainty when responding to complex technical questions. Instead, they tend to provide explanations that acknowledge context, limitations, and trade-offs. However, the way a prompt is written can strongly influence how confident the response appears. Direct prompts often produce answers that sound more decisive, while prompts that encourage caution lead to more explicit acknowledgment of uncertainty.