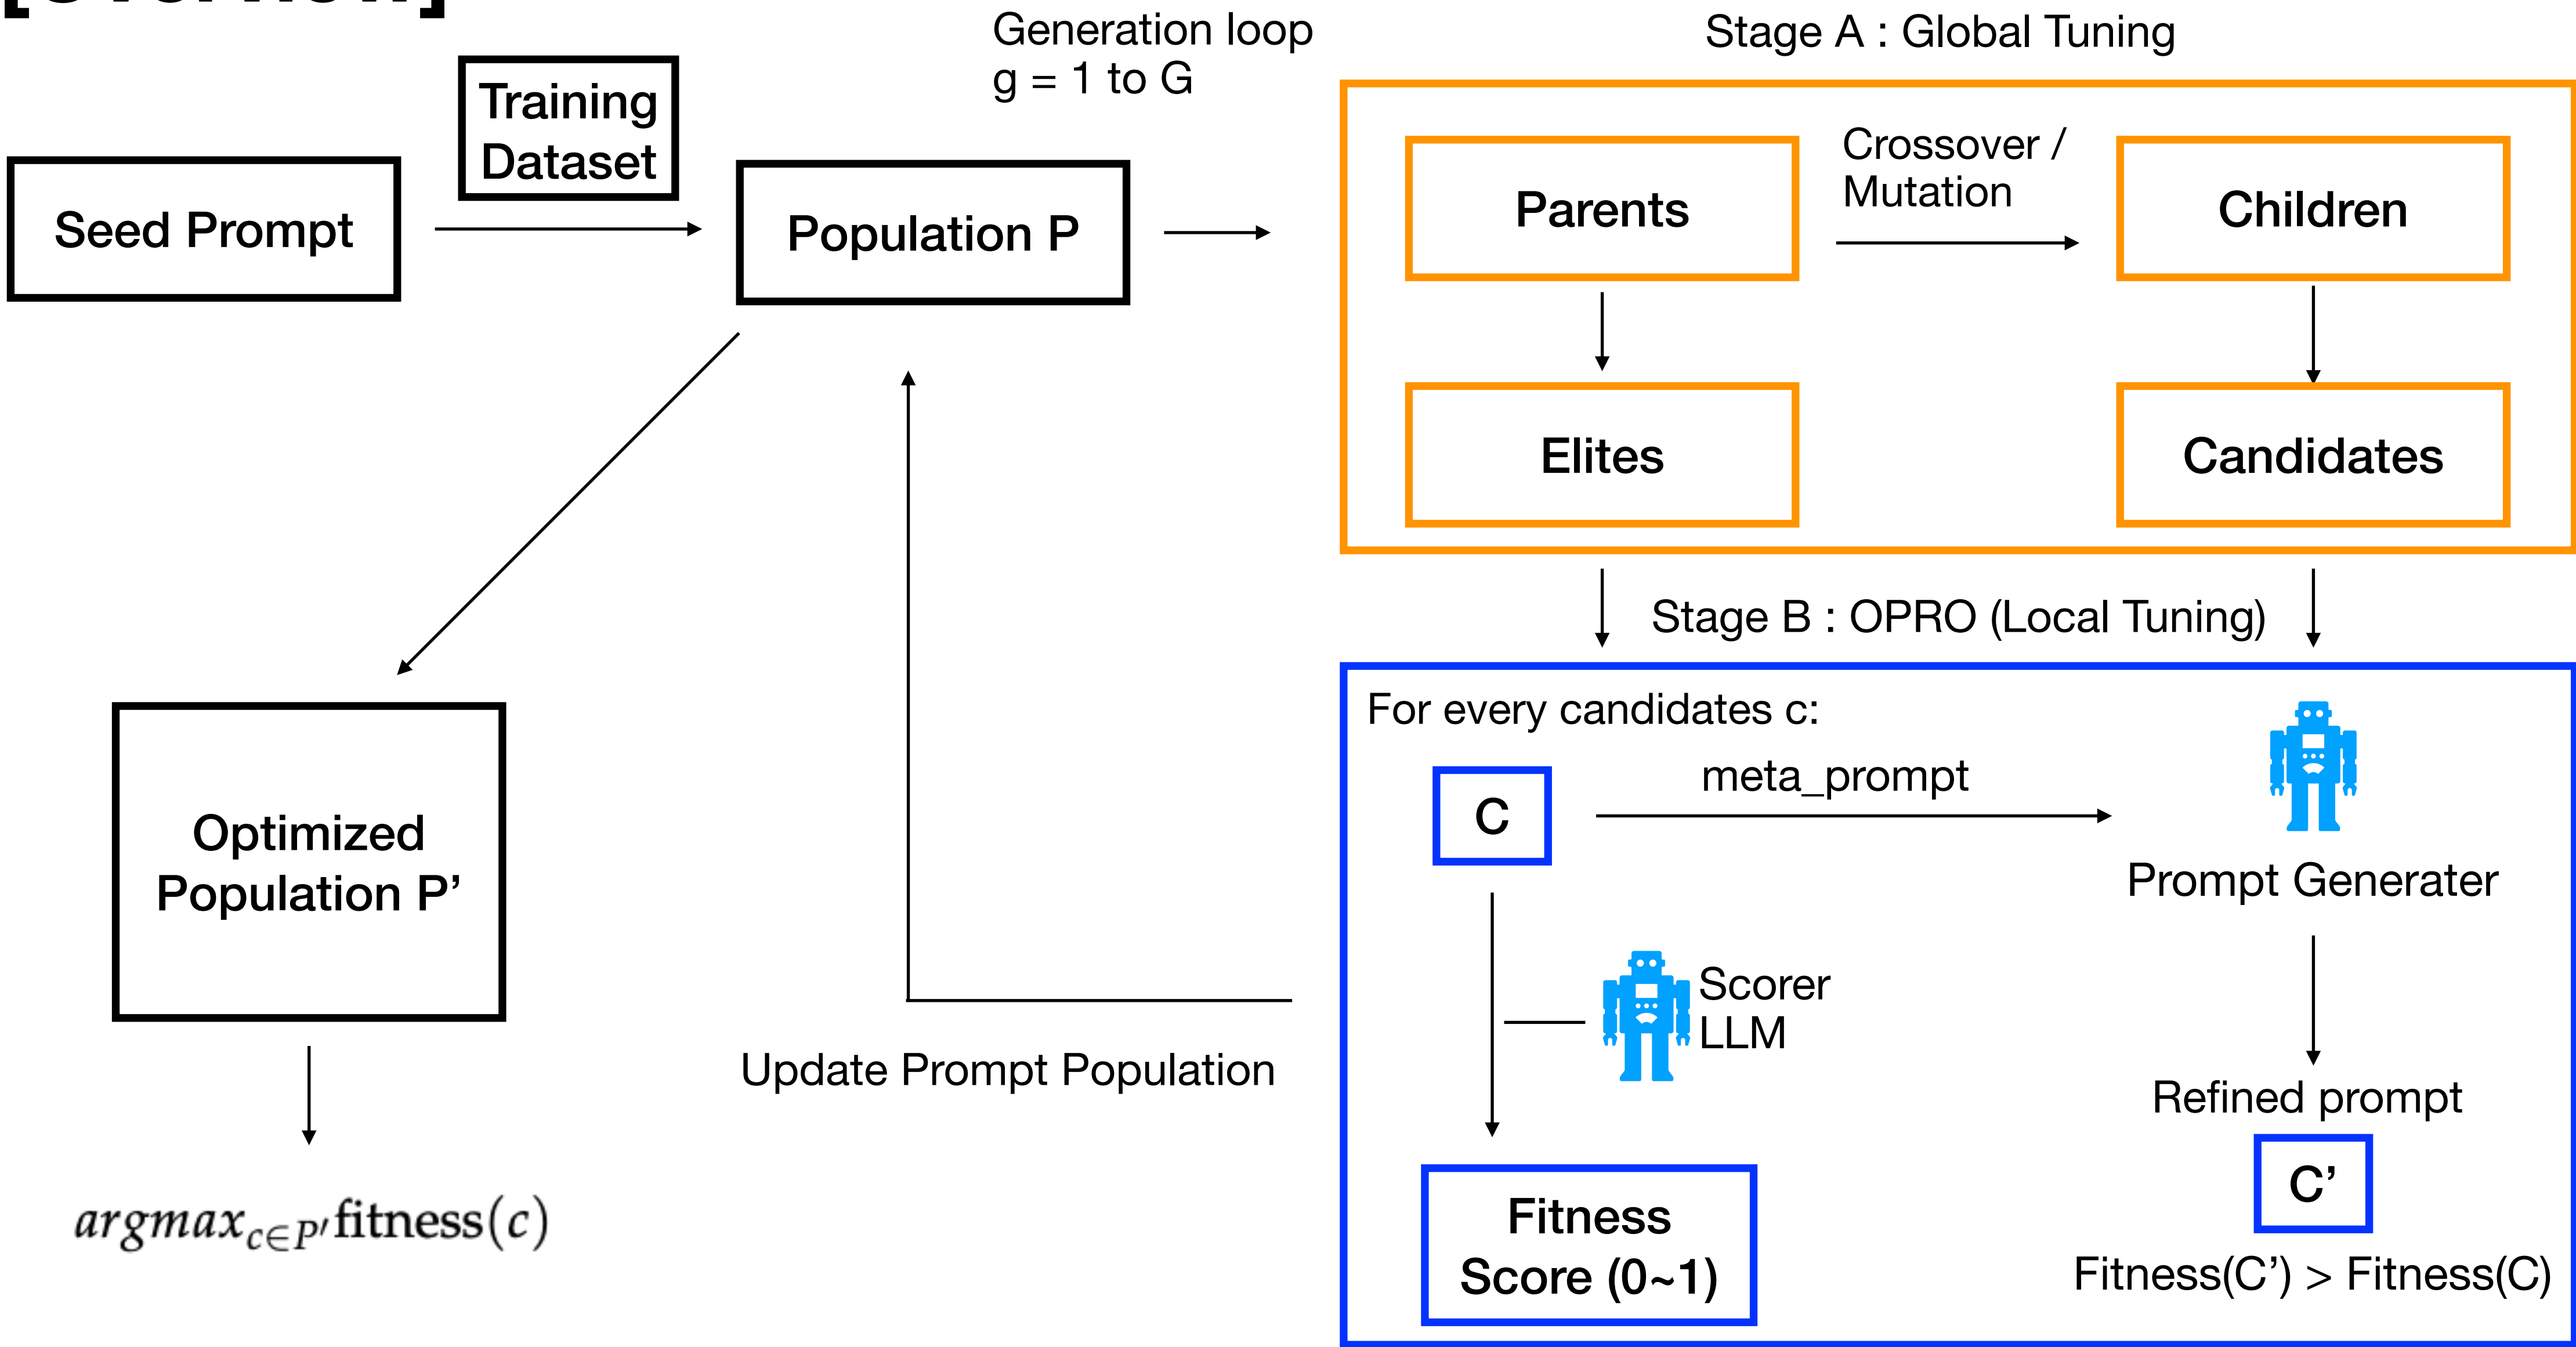


Prompt Optimization with EvoPrompt and OPRO

CAS4133 LARGE LANGUAGE MODEL
LLM_Final Project

Sunghyun PARK

[Overview]



0) Setup

- **LLM Scorer** : unsloth/Llama-3.1-8B-Instruct-unsloth-bnb-4bit or unsloth/Mistral-7B-Instruct-v0.2-bnb-4bit
- **LLM Prompt Generator** : unsloth/Llama-3.1-8B-Instruct-unsloth-bnb-4bit or unsloth/Mistral-7B-Instruct-v0.2-bnb-4bit
- **Dataset** : openai/gsm8k

1) Seed Prompt

Seed Prompt

- Solve the problem step-by-step in a short paragraph, then write the final result on a new line that starts with 'Answer:'.
- Think through the solution out loud in one or two sentences. Conclude with the exact result in the format: Answer: <result>.
- Provide a concise chain-of-thought (no more than three sentences) and finish with a single line beginning 'Answer:' followed by the solution.
- Explain your key reasoning briefly so the logic is clear, then end your response with 'Answer: <result>'.
- Reason internally but summarize the steps in one short paragraph. On the last line, state the solution prefixed by 'Answer:'
- Show a minimal reasoning trace (few sentences), ensuring the final line reads exactly: Answer: <final result>.

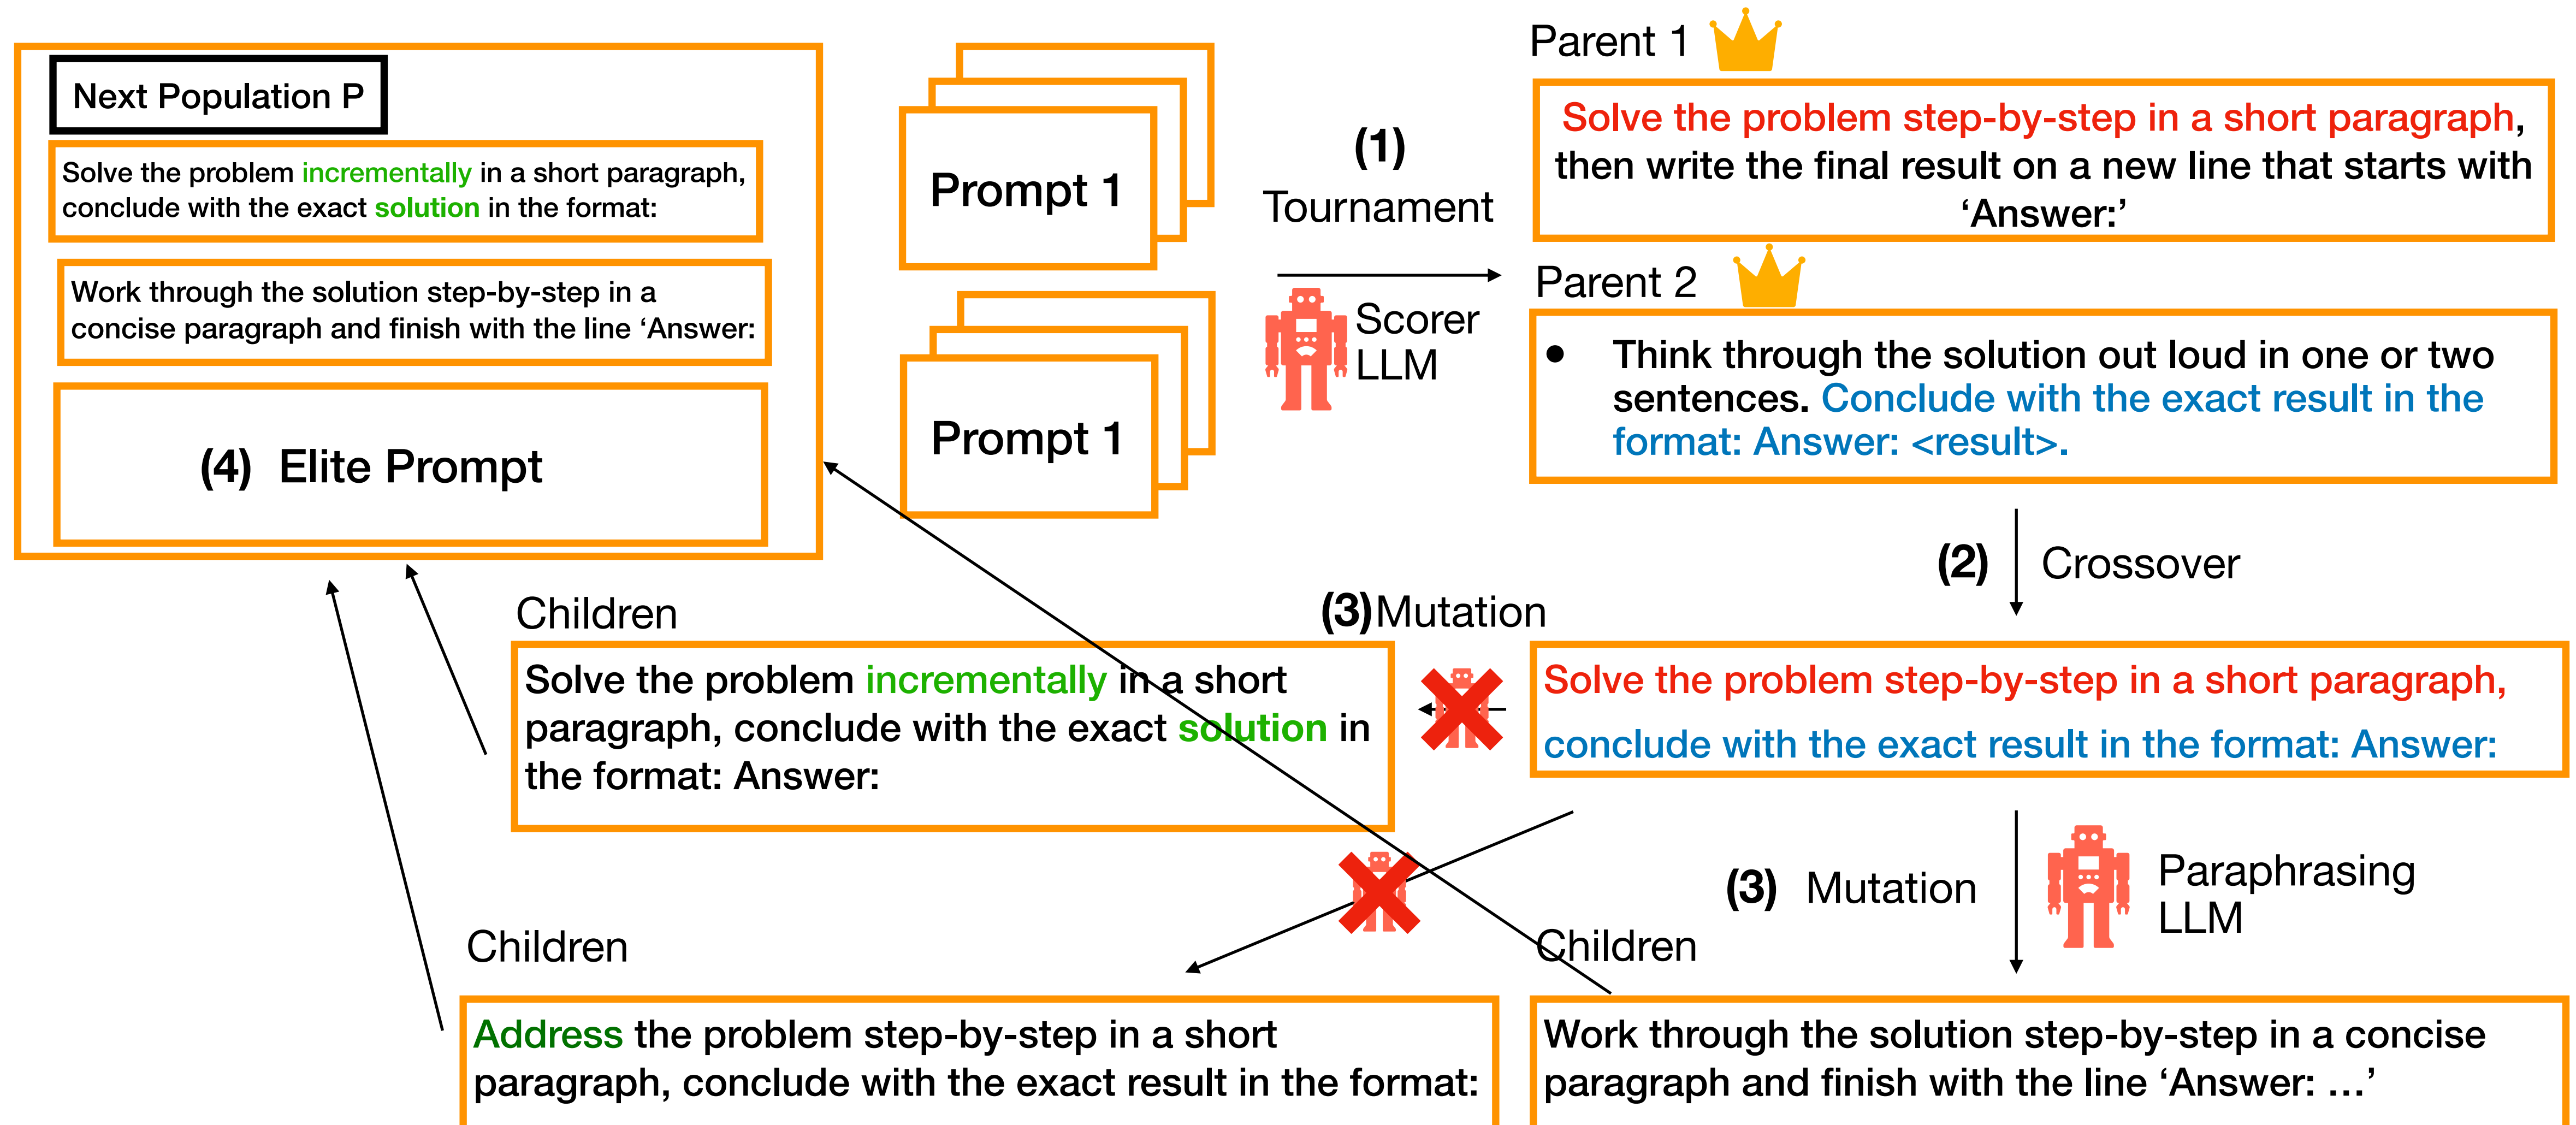
2) Stage A (Prompt Evolution Genetic Algorithm)

Global Search

- Purpose : Explore the prompt space globally before Stage B and provide diverse and high quality starting prompt to OPRO
- WorkFlow (per generation)
 - 1) **Parent Selection** - Tournament : Randomly pick 3 prompts and keep highest fitness winner and repeat P times to build the parent pool
 - 2) **Crossover** : Split each parent prompt at a random sentence boundary and swap the head of Parent A with the tail of Parent B -> new child prompt.
 - 3) **Mutation** : Ask LLM to paraphrase the child and if LLM fails, synonym swap on words
 - 4) **Elite Preservation** : Copy the top-E prompts unchanged into the next generation

2) Stage A (Prompt Evolution Genetic Algorithm)

Global Search



3) Stage B : OPRO

Local Prompt Refinement

- **Purpose** : Fine-tune each candidate prompt produced by stage A, and boost both answer accuracy and reasoning quality at a local level.

- **Workflow**

- 1) Build a meta-prompt

You are an expert prompt engineer. Your goal is to improve the provided prompt to increase its {task_objective}. Return ONLY the improved prompt without extra commentary.

Here are previous attempts and their scores:

Attempt {1}: '{p1_text}' (Score: {p1_score})

Attempt {2}: '{p2_text}' (Score: {p2_score})

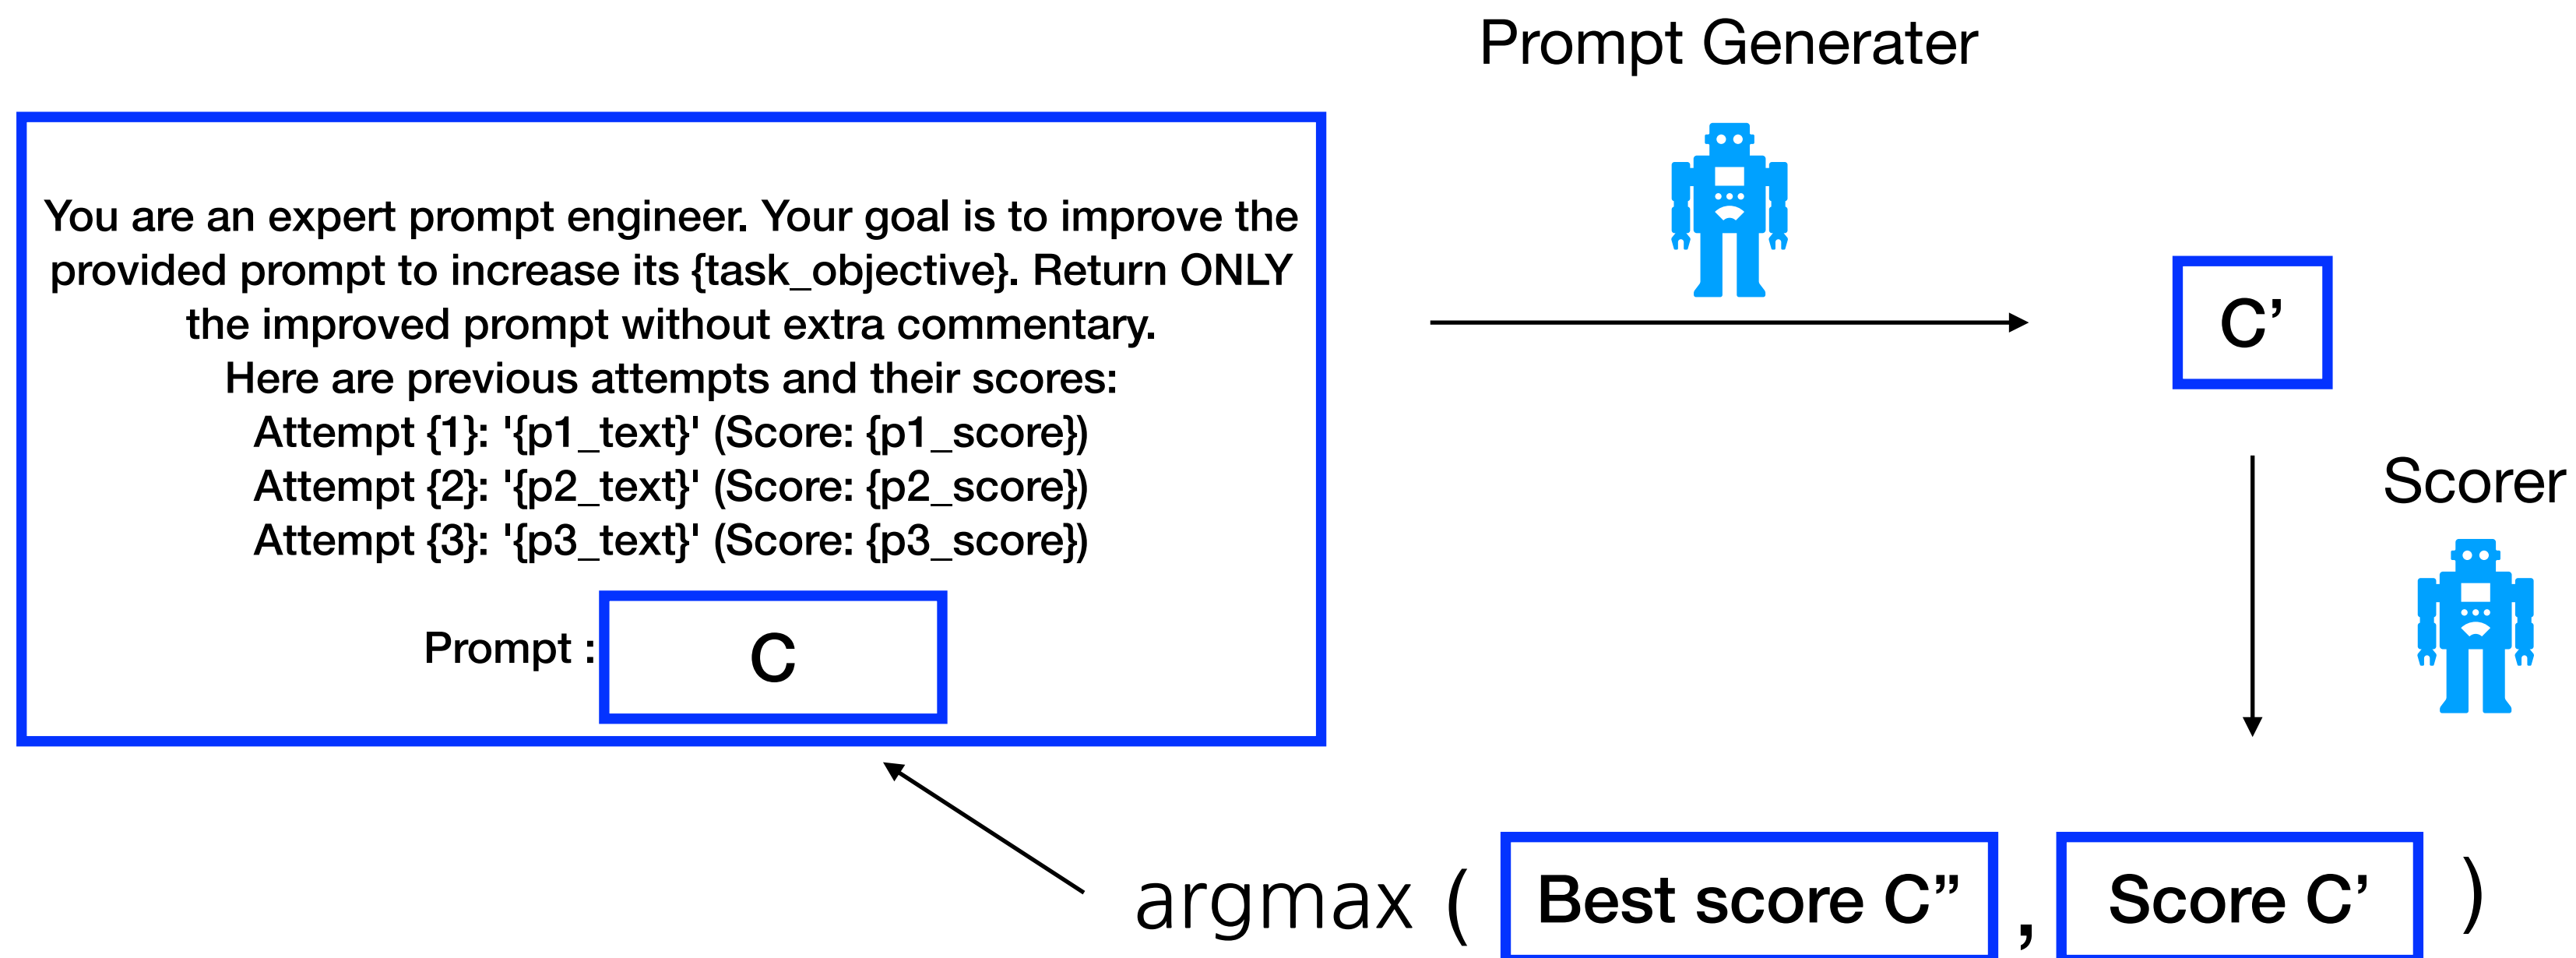
Attempt {3}: '{p3_text}' (Score: {p3_score})

Prompt : Current Prompt

3) Stage B : OPRO

Local Prompt Refinement

2) Generate Prompt and evaluate



(OPRO Terminates the loop after four consecutive non-improving steps)

Optimized Prompt

Optimized by LLaMA-8B

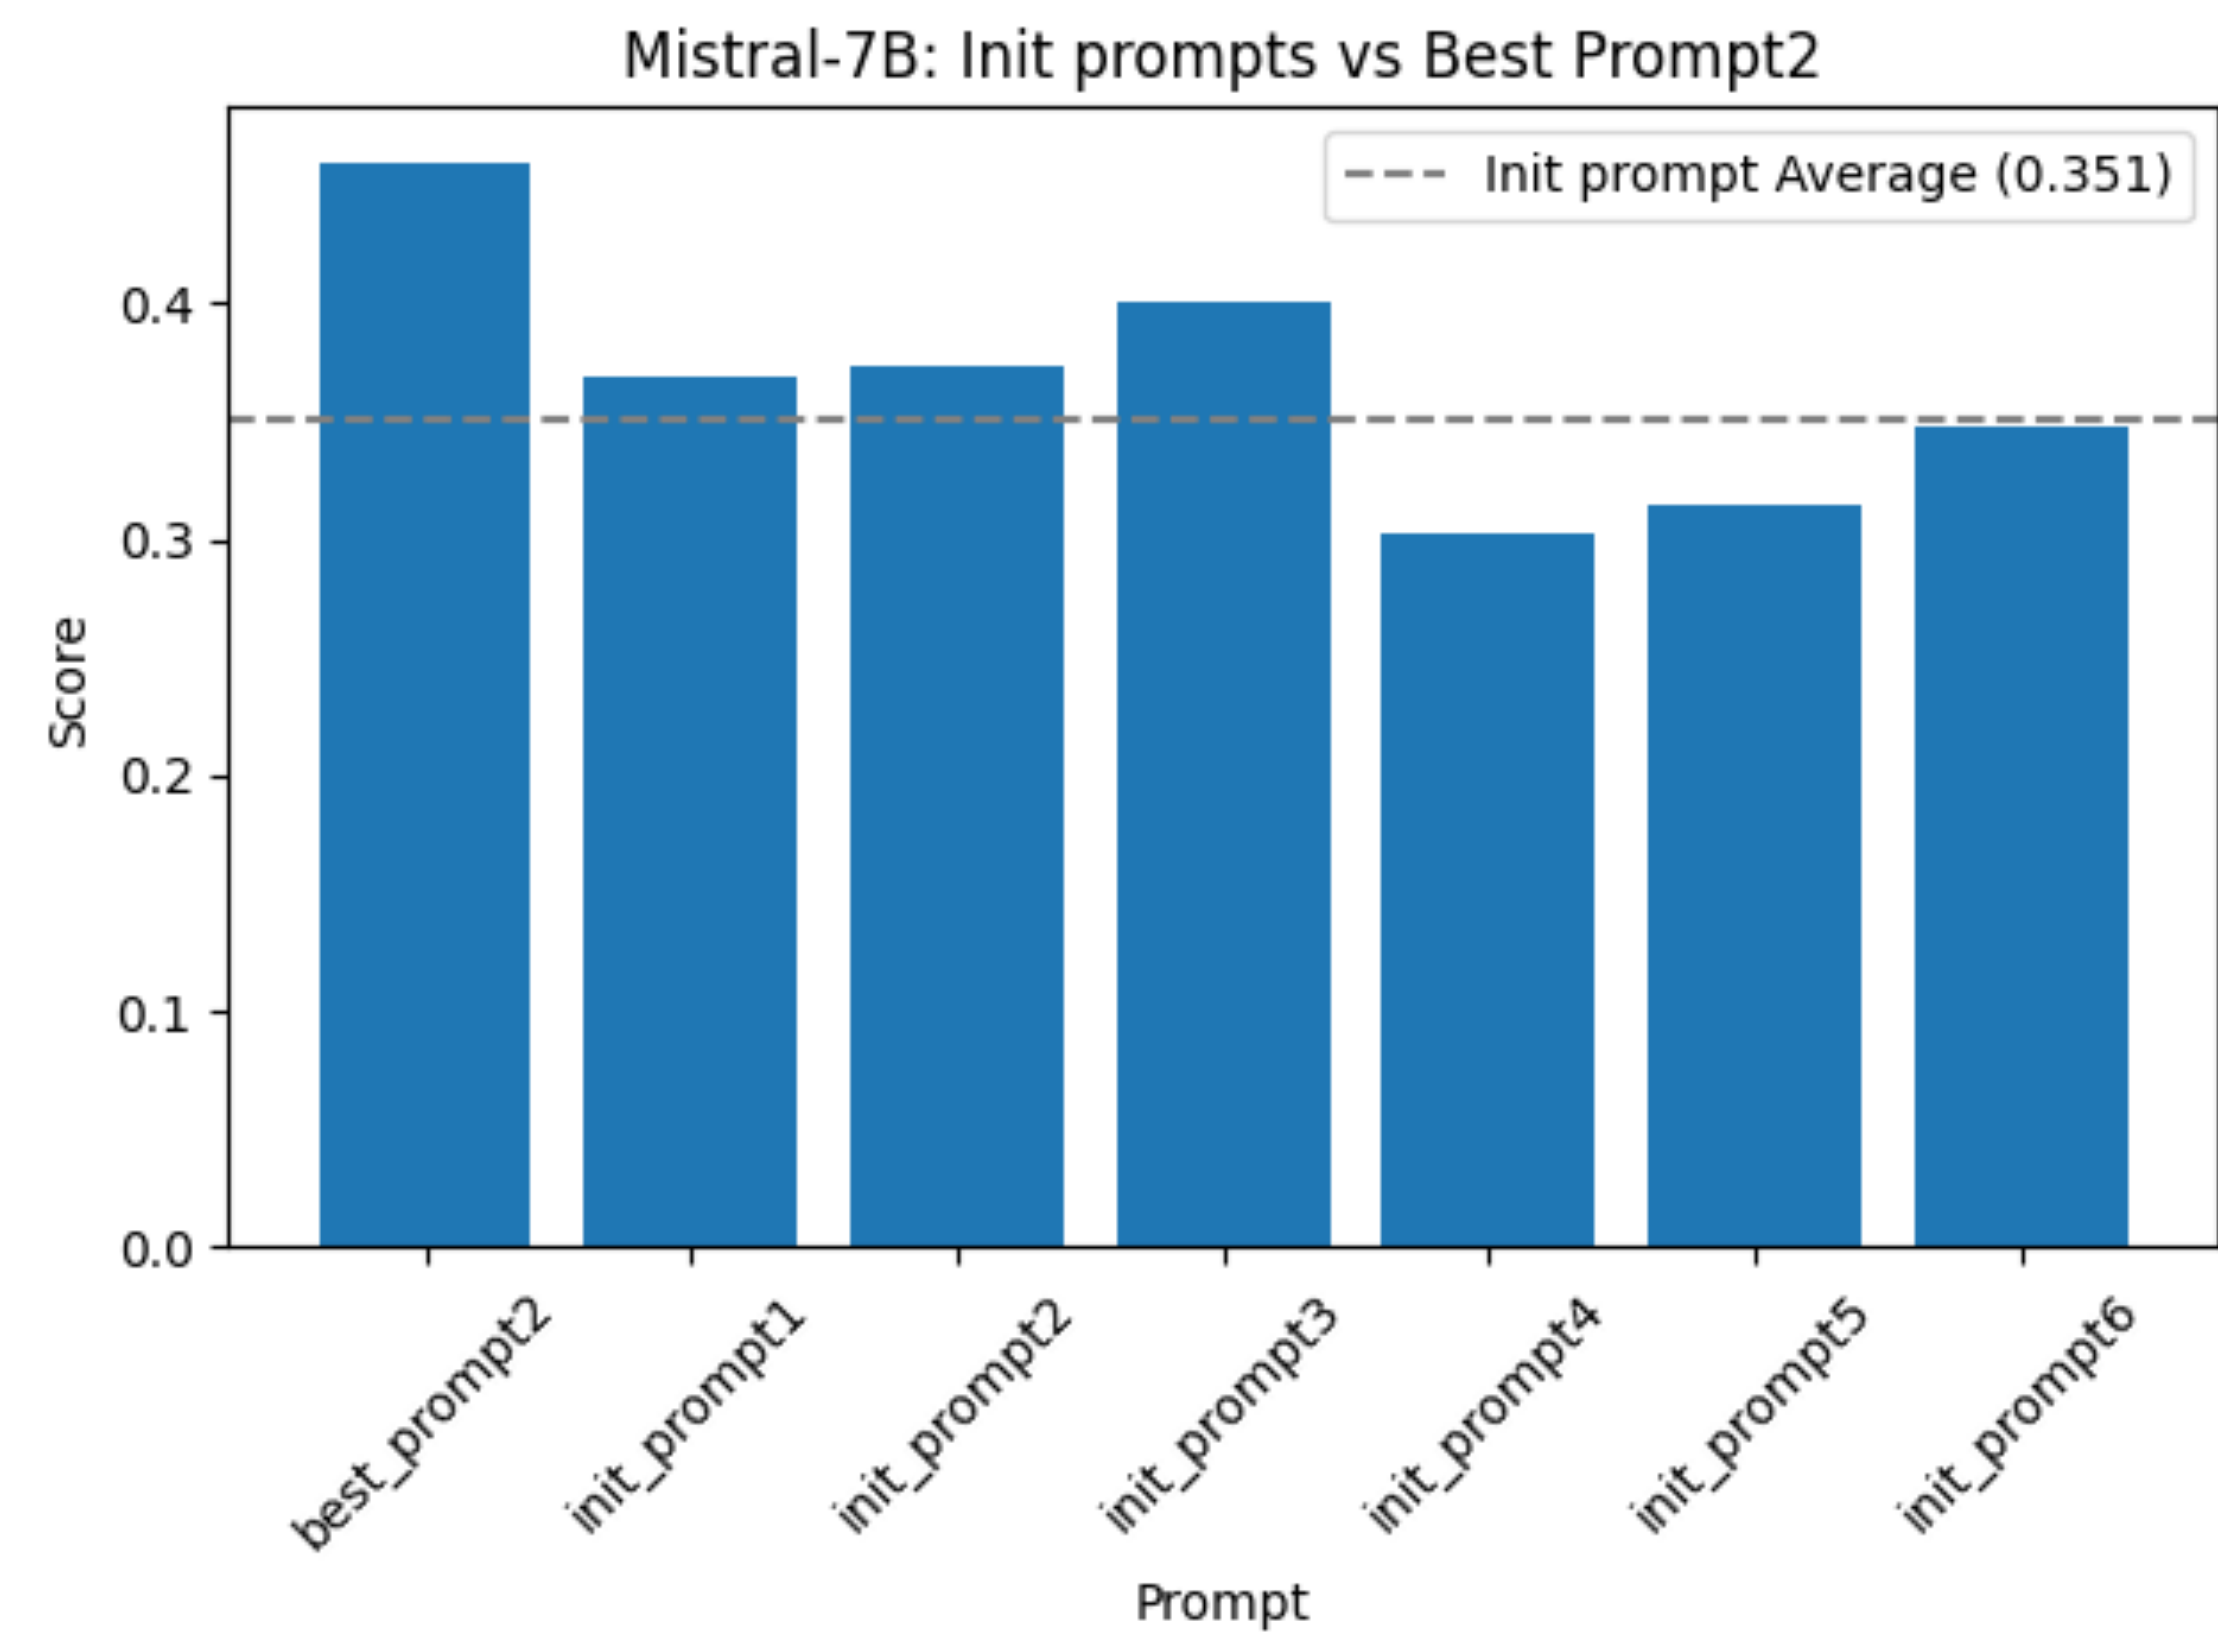
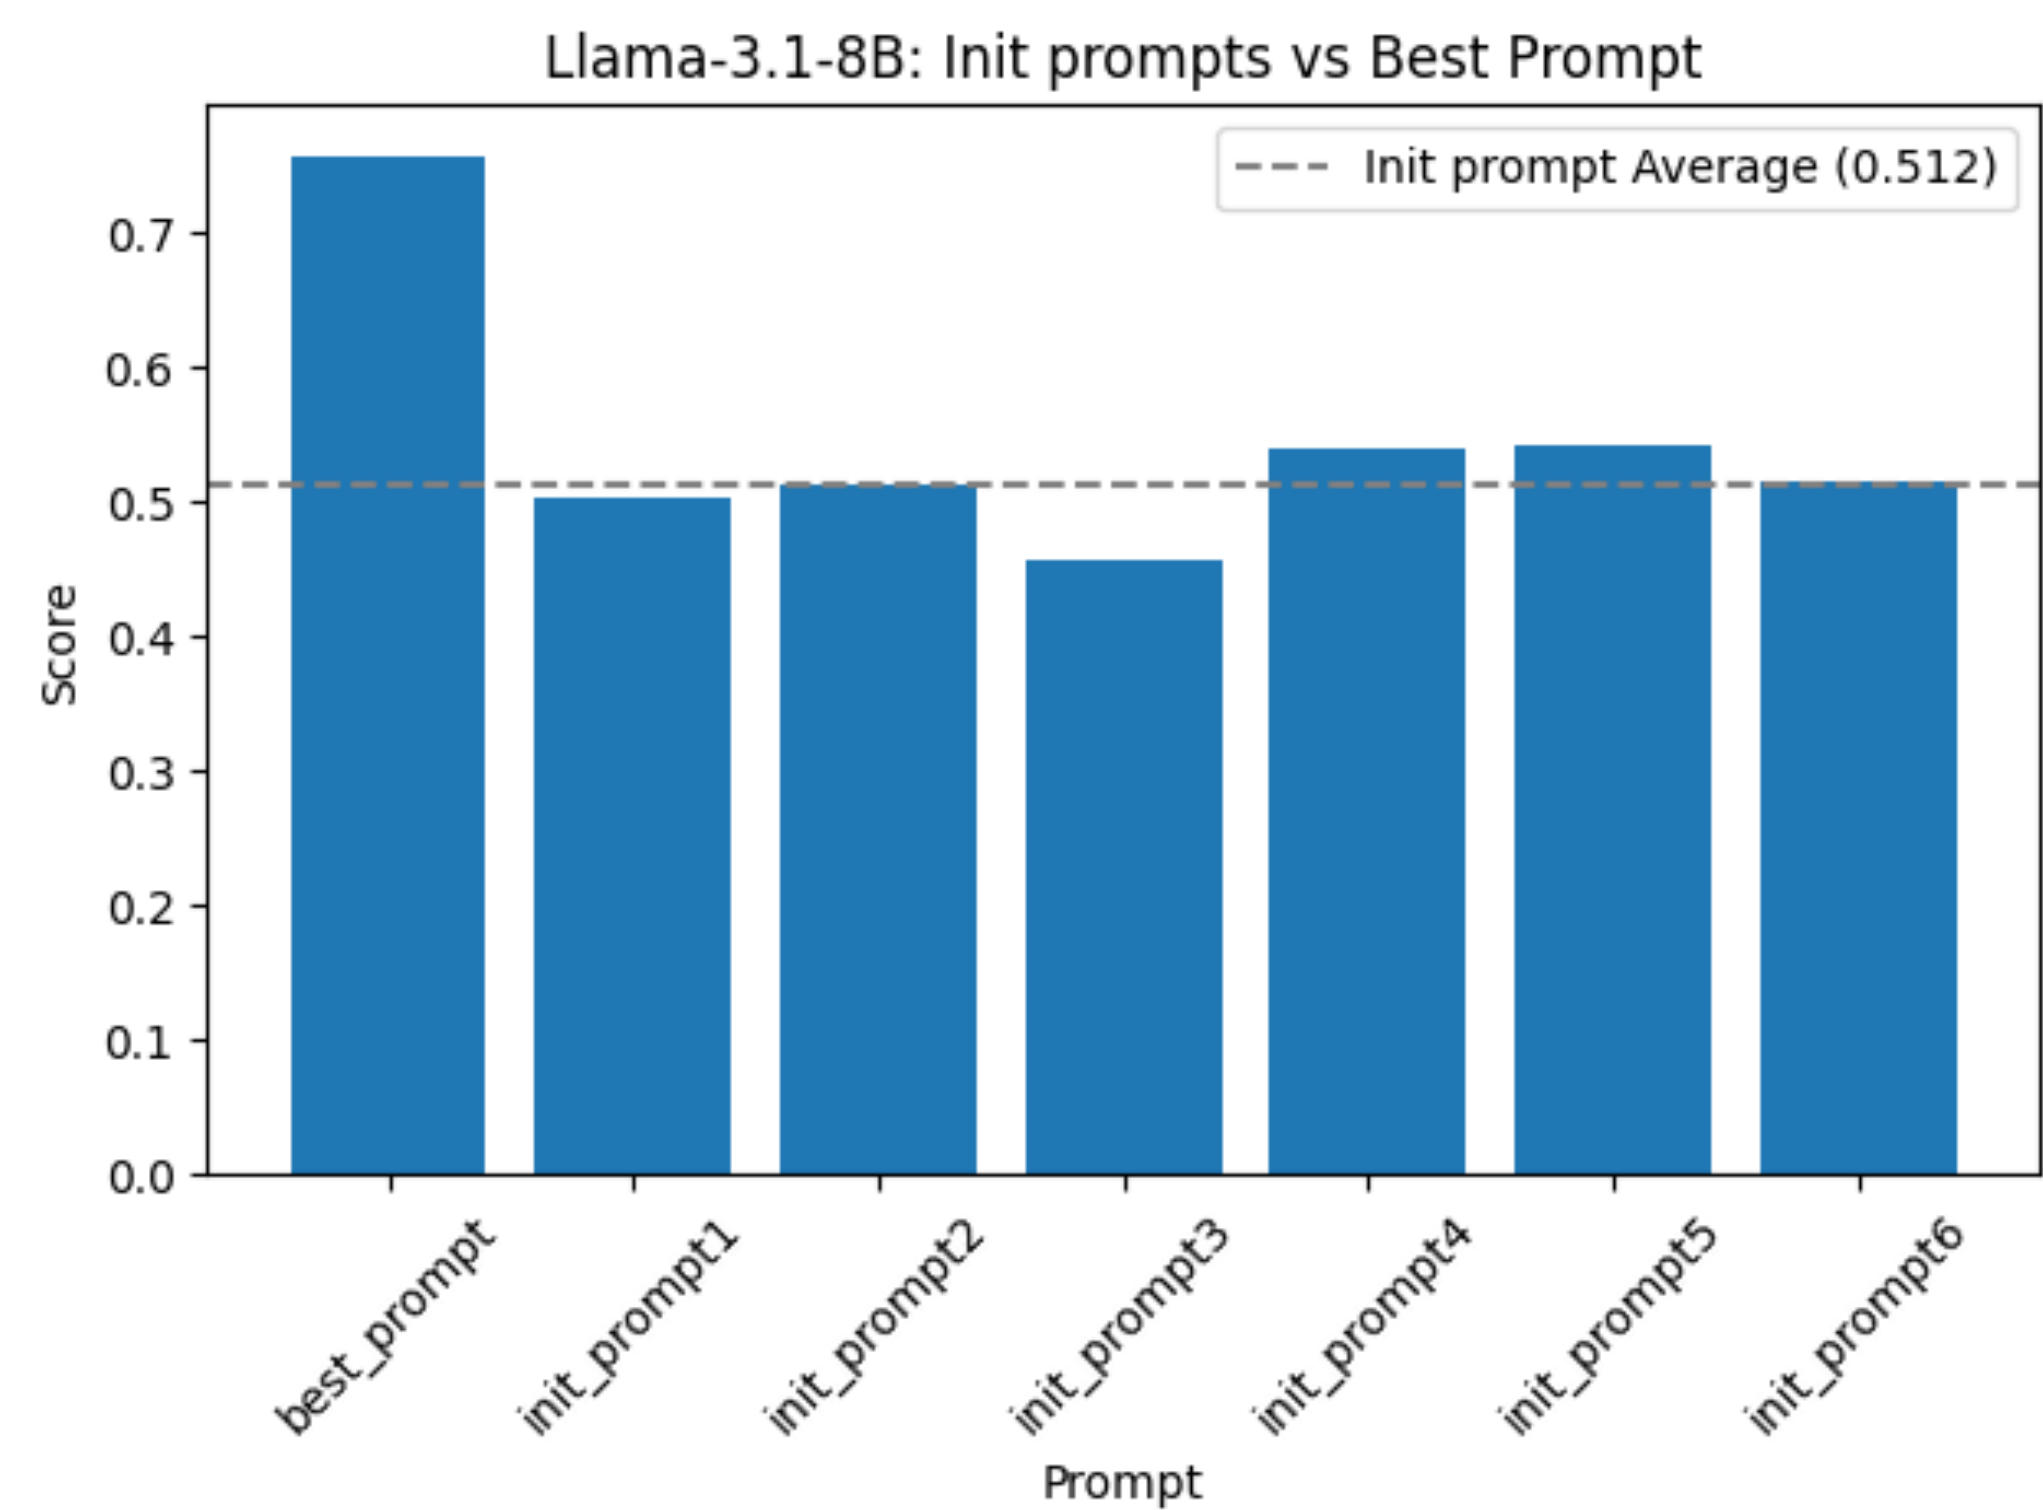
Describe the step-by-step reasoning process, including relevant mathematical equations, theorems, or logical deductions, that led to the final answer. Present the solution in a clear and concise format, labeled as 'Answer:', and provide necessary context or assumptions influencing the outcome. Explain underlying mathematical concepts or principles and any relevant theorems or equations used to arrive at the solution.

Optimized by Mistral-7B

PROMPT: Please provide a clear and concise explanation (maximum of three sentences) leading to the calculation of the answer. Begin your response with the word "Answer:" followed by the numerical solution.

Evaluation Results

GSM8K An evaluation was performed on the GSM8K test set.



Evaluation Results

Drawback : Since each prompt is optimized to a specific model, you must optimize it separately for each one. A prompt that is been fine tuned for one model may actually perform worse than the original initial prompt.

