How to Engineer a "Good" Prompt

A prompt for Generative AI is the statement guiding the AI to create contextually relevant outputs. A prompt influences the output's richness and accuracy. A good prompt should be **specific**, **precise**, **positive**, **ongoing**, and **experimental**.

There are many ways to engineer a prompt effectively. Some tools are provided below.

Type of Prompt	How it Works
Flipped Interaction	Ask the GenAI to ask you questions in order to create a more comprehensive response. *Example*: "Ask me about [insert here] in order to understand it better. When you have enough information, [insert task here]. Start by asking me the first question."
Reflect on Reflection (RoR)	Ask the model to evaluate the response that it generated and then to rewrite the response based on the evaluation. *Example*: "Evaluate the [completed task here] based on [parameters]. Now, redo the task."
Chain of Thought (CoT)	Ask the mode to show its work step-by-step or explain its reasoning. *Example*: "Show your work"; "Show how you arrived at this answer"; "Show the steps you used to arrive at this answer."
A Prompt Game	Create a game with the GenAI to practice effective prompting. *Example*: "I want you to play a game with me about [task here]. The purpose of this game is to improve my ability to prompt you. When I give you an unclear prompt, let me know how to improve it."
Cognitive Verifier	Ensure the LLM is clear on the query. It will ask questions to better understand and refine it. *Example*: "From now on, when I ask a question about [topic here], ask me questions based on [specific ideas around topic] to better understand my question. Then combine these into a response."
Few Shot	The model is trained on a few examples (usually a small number) per new class. *Example*: Provide more than one example; this is more practical than single shot (below) as the LLM has more information to work with.
Single Shot	A "single shot" prompt refers to a concise and self-contained input provided to a large language model in a single instance, typically consisting of a short sentence or a few words. *Example*: Contains a concise and self-contained input into the LLM – this does not rely on context from previous interaction.

Subdivide Questions	LLMs can often reason better if a question is subdivided into additional questions. *Example*: Provide multiple questions; ask the LLM to combine the responses.
The QEC Model	Composed of: Query (task)/Example (how)/Context (additional information) *Example*: Create an [output here]. Use the following criteria [provide context]. This output is intended/used for [provide context].
The Audience Persona	Ask the AI to produce an output for an audience persona. *Example*: Adopt [insert persona here] in answering this question.