Machine Learning 01

Kihyun Shin
DMSE, HBNU

TheMAD

Generative AI (Gen AI)

- The most important thing in dealting with Gen AI is Input and output.

- Output : expected from Gen AI

- Input : what kinds of <u>prompt</u> do you have to write to get proper

  Output **Input from user**

- To get good prompt: 5-step prompt frameworks

  - **Thoughtfully Create Really Excellent Input**

TheMAD

Design Great Prompts

- 5-step prompt frameworks

    - **T**ask (What to do-persona, including formats)

    - **C**ontext (specific situation)

    - **R**eferences (experience and record)

    - **E**valuate (evaluataion, true/hallucination)

    - **I**terate (repeat)

- **Thoughtfully Create Really Excellent Input**

TheMAD

TCREI: Specify the task

- Good prompt : starting from specifying the task thoroughly
- Think of it as asking a person for help rather than giving commands to a computer.
  - Write a list.
  - Write a speech
  - Create an image.
- When specifying a task, clearly and precisely describe what you want. If the task is ambiguous, the results may also be incorrect.
- To avoid ambiguity, you can utilize Persona or Format.
  - Persona: Assign a specific role to Gen AI, such as a scientist, an industry expert, or a filmmaker. You can also set the target audience to adjust the tone and complexity of the output. (For example, imagine the audience is students or company executives.)
  - Format: Control the output format. Instead of simple text, you can request structured outputs, receive data in Excel, or generate results in graphical form.
- Example
  - "You are a film critic specializing in Italian cinema. Create a table listing the best Italian films from the 1970s and classify them into thriller, drama, and comedy genres. Additionally, provide a 100-character summary for each film, including the director and release year."

TheMAD

TCREI: Include necessary context

- Including an excessive amount of details can be helpful in prompting.

- The more specific elements you include, the higher the quality of the output.

- Provide as much background information as possible (goal, purpose, previous attempts).

- These detailed elements are called context.

- A prompt can be written simply as follows:

  - **How was DNA discovered?**

- If you include detailed elements and background information:

  - **"You are a professor giving a lecture on DNA at a university. Tell me a few interesting events related to the discovery of DNA, and explain how this discovery has impacted the world. However, make sure to explain it in a way that is easy for the general public, who may not be familiar with science, to understand. You have previously received feedback from students that this lecture is boring and difficult. Therefore, make it as engaging and easy to understand as possible."**

TheMAD

TCREI: Provide references

- References refer to examples or similar past results that Gen AI can analyze and use as a reference.

- For example, when visiting a hair designer, showing a picture of the desired hairstyle is much faster and more effective than trying to describe it with words.

- The type of files that can be uploaded depends on the Gen AI tool being used. You can upload text, images, audio, video, code, Excel, PPT, etc., to clarify your request.

- However, providing too many references may lead to a one-sided result. 2-5 references are usually sufficient, but this depends on the situation.

- Example Prompt:

    - "Write a product description for a wristwatch as if it were published in a magazine."

- By adding references, you can refine your request further:

    - "Write a one-paragraph product description. The product is a high-performance wristwatch, resistant to scratches, and waterproof up to 30 meters. Base the writing style on the following examples."
        - Sunglasses: Our latest collection of handcrafted, heritage-inspired sunglasses features details like UV-protective lenses in shades specifically chosen to complement each frame   all at a price that won't break the bank. Plus, we made these sunglasses with vintage-inspired acetate frames and a keyhole bridge.

        - Cardholder: Crafted in smooth Italian leather, this double-sided cardholder is designed to carry your cash and credit cards without the bulk of a full wallet. Fun fact: this cardholder is made in Naples, Italy, and will look great when you treat your friends to a round of summer spritzes."

TheMAD

TCREI: Evaluate your output

- Each Gen AI model is trained using different datasets and techniques.

- Therefore, different models may excel in different areas   some may be better at coding, while others are better at brainstorming, etc.

- However, regardless of the model, even when using the same prompt, the results may vary slightly.

- Thus, evaluating the output is very important.

- Before using AI-generated results, you should check whether:

- The result is accurate. There is any bias. The information is relevant.

- It aligns with previous findings.

- If the output is not as desired, you need to iterate and refine your prompt.

Especially, Gen AI results are highly prone to fabrication (so called hallucinations), so a strict fact-checking process is essential.

TheMAD

TCREI: Take an iterative approach

- In most cases, the first prompt will not produce the desired result.
- Always follow ABI (Always Be Iterating): Continuously refine and adjust your prompt while repeating the process.
- Always keep in mind the five key principles mentioned earlier and keep refining your prompt until you achieve the desired outcome.

  - 5-step prompt frameworks

    - **T**ask (what to do-persona, including formats)

    - **C**ontext (specific situations)

    - **R**eferences (experience/record)

    - **E**valuate (evaluation, true/hallucinations)

    - **I**terate (repeat)

  - **Thoughtfully Create Really Excellent Input**

TheMAD

Better output through iteration

- **Iteration process**

- **example**

    - **What is the latest technology related to the restaurant industry?**

        - **This question is very broad, which may result in a vague or overly general response.**

        - **It is necessary to iterate and refine the prompt until you obtain the desired level of detail.**

        - **There are four steps in this process.**

TheMAD

Better output through iteration

- **Revisiting the 5-Step Prompt Framework**
  - Create a bulleted list of the latest technologies related to the restaurant industry. Focus on urban areas and limit the results to places that use only locally sourced ingredients.
- **Breaking Down the Prompt into Steps**
  - 1. Create a bulleted list of the latest technologies in the restaurant industry.
  - 2. Summarize only the technologies that have had a significant impact on urban restaurants.
  - 3. Analyze the advantages and disadvantages when these technologies are used in restaurants that rely solely on locally sourced ingredients.
- **Adding Specific Conditions (Category, Length, Format, etc.)**
  - Create a bulleted list of the latest restaurant-related technologies.
  - Search specifically in urban areas and limit the results to restaurants using only local ingredients.
  - Additional conditions: The city must have a population of at least 500,000 and only show technologies related to vegetarian cuisine.
- **Modifying the Prompt or Changing the Context**
  - "I am starting a restaurant business. I plan to use only ingredients sourced within 50 km of the restaurant. You are a food enthusiast living in a nearby metropolitan area. As a gourmet, please create a list of things I should consider before opening my restaurant."

TheMAD

Understand multimodal prompting

- If you were a painter, creating a masterpiece with just one color would be difficult but possible.

- However, using a variety of colors would make painting much easier.

- Multimodal Prompting refers to inputting multiple types of data into a prompt simultaneously. For example : Text + Image, Text + Excel, Text + Graphs, etc.

- When it is difficult to describe exactly what we want using only text, multimodal prompting helps us effectively communicate our intent to the AI.

- In reality, we already use multimodal methods extensively.

- Think about giving a presentation: **(To convey our message more effectively to the audience.)**
  - We use PowerPoint slides (images).
  - We sometimes include videos (visual aid).
  - We speak while presenting (audio).
  - We highlight text within the slides.
  - We make eye contact to engage the audience (interaction).
  - We adjust our tone of voice for emphasis.

TheMAD

Understand multimodal prompting

- Multimodal prompting is one of the most effective ways to transform different types of information.

- Examples of Multimodal Prompting:
  - Requesting a recipe by showing a photo of the inside of a refrigerator.
  - Generating a digital teaser for a social media collaboration between two brands by providing their logos and color schemes.
  - Creating a short story about a forest and describing its atmosphere by uploading recorded sounds and pictures and asking the AI for an interpretation.
  - Uploading lecture materials and requesting a summary and midterm exam questions based on the content.
  - Providing a conference program and asking for a summary with extracted key sessions relevant to the user's interests, formatted into a table.

- Just like any other AI interaction, multimodal prompting follows the same 5-step prompt framework discussed earlier.

TheMAD

Prompt Responsibly

- Gen AI is not flawless   it can: Make mistakes. Produce biased or stereotyped results. Misinterpret information. Contain inappropriate content.
- Therefore, human-in-the-loop (HITL) is essential   humans must be involved in the evaluation process.

- **Biases and stereotypes**
    - Gen AI can generate biased results based on region, religion, race, age, gender, and historical perspectives.
    - **Specify diversity**
        - Instead of asking for an image of "a table with delicious food," specify "a table with delicious food from around the world" or mention a specific region or cuisine to get a more accurate and diverse result.
    - **Challenge assumptions**
        - If Gen AI generates an image of "a table with delicious food" that mostly features Southeast Asian cuisine, you can point out this bias and ask for corrections.
        - By continuously refining and iterating prompts, AI models can learn to provide better and more inclusive outputs over time.

TheMAD

Prompt Responsibly

- **Hallucinations**
  - When Gen AI provides false or misleading information, it is referred to as hallucination.
  - This can happen at any time, but it is especially common when: The prompt is vague or ambiguous. The AI lacks sufficient data to generate an accurate response.
  - To minimize hallucinations, follow these steps:
    - **Fact-check and cross-reference**
      - Some Gen AI models integrate search engines to verify the accuracy of their responses.
      - If the topic requires expertise, consult a subject matter expert.
      - Another approach is to ask the same question to multiple Gen AI models and compare their responses for consistency.

    - **Use clearer or more detailed language**
      - If the prompt is unclear, Gen AI may misinterpret it and provide an incorrect response. Example: Asking "Why is Hanyang the capital of South Korea?"
      - Some AI models might assume this statement is true and generate a fictional explanation instead of correcting the assumption.
      - To avoid this, make sure the prompt is clear and precise.

TheMAD

Prompt Responsibly

- **Inconsistencies and relevance issues**
    - Gen AI sometimes fails to understand context properly.
    - A common example is with idiomatic expressions:
    - If you ask AI for a "response that breaks the mold," instead of giving a creative answer, it might literally explain how to break a physical mold or window.
    - To resolve these issues, use the following strategies:

        - Provide references
            - Giving specific references helps AI understand the intended context and reduces the chances of generating irrelevant responses.

        - Provide more context
            - More detailed context improves AI comprehension.
            - Instead of saying "Give me an answer that breaks the mold,"
            - A better prompt would be: "Provide a unique and creative response that challenges conventional thinking."
            - This clarifies the expectation, making it easier for AI to generate the desired output.

TheMAD

Checklist for using Gen AI responsibly at work

- **Consider the effects of using AI**
  - Although AI may seem capable of solving all problems, that is not always the case.
  - Before using AI, ask yourself: Is AI the right tool for solving this problem? Could AI-generated results be biased or cause harm to certain groups?

- **Get approval to use AI**
  - Some organizations prohibit or restrict AI usage based on company policies.
  - Always check company guidelines before using AI tools.

- **Consider the privacy and security implications**
  - Be mindful of personal data and confidential company information to prevent leaks.

- **Evaluate all AI-generated content**
  - Every AI-generated result should go through human review before being used or published.

- **Disclose your use of AI**
  - If AI was used in generating content, transparency is necessary   it should be explicitly stated.

TheMAD

Tone and style

- You may have noticed that Gen AI sometimes responds like a robot and at other times in a friendly manner.

- This is because Gen AI can adjust its tone and style based on the given prompt.

- **Contextual understanding**
  - Gen AI is trained on vast amounts of conversational data, allowing it to understand context.
  - For example, if you ask: "Have you been eating well?" Instead of responding mechanically, the AI might say: "Thanks for asking! I've been eating well. Have you been busy lately?" This demonstrates AI's ability to grasp conversational context and generate natural responses.

  - How Gen AI Understands and Adapts Tone & Style
    - Determines the identity of the user (based on conversation style and context).
    - Recognizes the structure of the question and sentence endings to infer tone.
    - Learns language patterns from references (if provided).
    - Adapts to the user's preferences over time.
    - Maintains consistency in the chosen language and style.
    - Processes input and tailors responses to match the requested tone and style.

TheMAD

Tone and style

- **Strategies for tone and style in prompting**
    - Sometimes, you may need a specific tone and style when writing an email, presentation, or article.
    - While Gen AI generally adjusts tone and style automatically, it may not always be perfect.
    - For example,
        - "Make it academic" try "Write like a professor with depth and scholarly insight."
        - "Make it persuasive" try "Write in a way that grabs attention and builds trust with the audience."
        - "Make it satirical" try "Write with dry humor, like a stand-up comedian."
        - "Make it inspiring" try "Write in a way that is motivational and uplifting."
        - "Make it simple" try "Explain it as if you were talking to a kindergarten student."

    - Providing clear and detailed instructions about tone and style helps ensure better results.
    - Specify certain words to use or avoid (e.g., avoid jargon, use casual expressions).
    - Provide a reference with the desired tone & style (e.g., a sample text that AI can follow).
    - Use iteration (ABI - Always Be Iterating) to refine the prompt until you get the perfect result.

TheMAD

Summarizing documents

- Gen AI can be used to summarize documents, meeting minutes, and research papers efficiently.

- In the early days, AI models had limited token capacity, making it difficult to process large documents at once. However, modern AI models can now summarize significantly larger documents in a single request.

- Token

  - A token is the smallest unit that AI uses to process language.

  - A token can be a single character, part of a word, or multiple words combined.

  - When a user enters a prompt, AI breaks the sentence into tokens for processing.

  - Early AI models could handle a few thousand tokens, while recent models can process over a million tokens.

- If a document exceeds the token limit, AI can still summarize it by:

- Breaking the document into smaller sections and summarizing each separately.

- Using a step-by-step summarization approach instead of condensing everything into one sentence.

- This prevents loss of important details and reduces hallucinations (fabricated information).

TheMAD

Strategies for data analysis

- If you have large datasets, Gen AI can be effectively used to:

- Analyze patterns, Identify correlations between data points, Detect trends and insights

- However, different Gen AI models have different capabilities.

- It's important to know what your model can and cannot do before performing analysis.

  - Text Analysis
    - Extract key points and summarize content from large text datasets.
    - Analyze tone (e.g., identifying sentiment in customer reviews).
    - Process thousands of reviews to develop business strategies based on trends.

  - Data Expansion
    - AI can generate additional synthetic data to make analysis easier.
    - However, it's crucial to distinguish between real data and AI-generated data.
    - While useful for observing overall trends, it may lack high reliability.

  - Question & Answering with Data
    - You can upload datasets and ask AI to explain or summarize key findings.
    - AI can quickly identify data distributions and trends, providing essential insights.

TheMAD

Strategies for data analysis

- Scenario Analysis
    - AI can examine past data and predict possible future outcomes.
    - Useful for market forecasting, financial trends, and business strategy planning.

- Image Analysis
    - AI can extract and interpret data from uploaded graphs and charts.
    - Example: Analyzing two years of sales records to identify trends and forecast future performance.

- Customer & Market Analysis
    - AI can analyze social media trends, charts, and graphs to determine:
    - What customers want and their purchasing behavior.
    - How the market is evolving and potential shifts in consumer demand.

1. **AI is excellent at uncovering hidden patterns in data.**
2. **Instead of expecting AI to conduct a full analysis,**
    → **Use open-ended questions to explore new insights.**
3. **Rather than asking AI "Analyze this data",**
    → **Ask "What are some possible ways to analyze this data?"**

TheMAD

Temperature

- You may have noticed that even when using the same prompt, Gen AI sometimes produces different responses.

- This happens because AI generates answers based on probability, selecting words based on likelihood rather than fixed rules.

- A well-structured prompt can reduce randomness and make responses more consistent.

- A vague prompt may lead to high variability and creativity.

- A precise prompt can minimize hallucinations and generate fact-based responses.

TheMAD

Temperature

- Gen AI does not select words completely at random when generating responses. Instead, it uses a sampling parameter to adjust the range of possible word choices. This parameter is called temperature.

- When temperature is high (e.g., 1.0), the range of word selection widens, allowing for more creative and diverse responses.

- When temperature is low (e.g., 0.1), the range of word selection narrows, resulting in simpler and more predictable responses.

- Most Gen AI models allow users to adjust this setting, or it can be controlled through prompt design.

"This morning, I went to the…"

Token	Probability of selection
🛒 Grocery store	.41
🧺 Laundromat	.26
🎞 Movies	.18
🩹 Chiropractor	.11
🏚 Haunted house	.02
⛰ Volcano	.019
🌙 Moon	.001

TheMAD

Top-k sampling

- Gen AI uses a sampling method based on how frequently certain words appear in generated results (total number of tokens).

- Increasing the Top-k value (e.g., 1) (similar to lowering temperature to 0.1)
  - Produces more consistent results.
  - Minimizes hallucinations.
  - Using the same prompt will always generate the same response.

- Decreasing the Top-k value (e.g., 0.1)
  - Produces more unique and diverse responses.
  - Ideal for generating creative and unpredictable outputs.

TheMAD

Top-p sampling

- Top-k sampling selects words based on the number of tokens, whereas Top-p sampling further refines the selection based on probability within the Top-k results.
- For example, imagine selecting 10 romance novels from a bookshelf (Top-k sampling).
    - If Top-p score is set to 0, only the most popular bestseller (the most frequently chosen book) will be selected.
    - Since bestsellers are widely known, there is a high chance you have already read it.
    - However, by setting the Top-p score to 0.8, the selection becomes more diverse, increasing the likelihood of discovering new or lesser-known books.
- Typically, Top-p is set to 1.0, meaning it does not affect sampling and allows for a more natural selection process.

TheMAD